Equilibrium Effects of Pay Transparency*

Zoe B. Cullen† Bobak Pakzad-Hurson‡

June 23, 2017

Click For Latest Version

Abstract

The public conversation about increasing pay transparency largely ignores equilibrium effects, namely how it leads firms to change hiring and wage-setting policies and workers to adjust bargaining strategies. In this paper, we study these effects with a methodologically diverse approach. Our analysis combines longitudinal study of thousands of workers and employers facing different levels of pay transparency on TaskRabbit, an online labor market, with a parsimonious equilibrium model of dynamic wage setting and negotiation. We find, theoretically and empirically, that increasing pay transparency can increase employment, decrease inequality in earnings, and shift surplus away from workers and toward their employer. Intermediate levels of pay transparency, achieved through a permissive environment to discuss relative pay, can exacerbate the gender pay gap by virtue of network effects. There may be a direct need for government intervention in order to maintain a desirable level of transparency. Any scheme in which employers vary transparency based on private characteristics is unsustainable, as the signal sent to prospective workers is sufficiently strong to cause unraveling toward full transparency. We observe this unraveling on TaskRabbit. We also conduct a field experiment on internet workers to investigate an alternative model in which wage compression is driven by social aversion to observed wage inequality. Our findings are consistent with our bargaining model but not with this alternative.

*We are very grateful for guidance from Susan Athey, Nick Bloom, Matt Jackson, Fuhito Kojima, Ed Lazear, Luigi Pistaferri, and Al Roth. We are indebted to James Flynn for technical support. We also thank Mohammad Akbarpour, Jose Maria Barrero, Doug Bernheim, Eric Budish, Gabriel Carrol, Arun Chandrasekhar, Isa Chaves, Piotr Dworczak, Jack Fanning, Chiara Farronato, Bob Hall, Gregor Jarosch, Scott Kominers, Maciej Kotowski, Jon Levin, Shengwu Li, Erik Madsen, Davide Malacrino, Alejandro Martinez, Paul Milgrom, Muriel Niederle, Ilya Segal, Isaac Sorkin, Jesse Shapiro, Takuo Sugaya, Emmanuel Vespa, Alistair Wilson, and various seminar attendees for helpful comments and suggestions. We appreciate TaskRabbit for providing essential data access. Pakzad-Hurson acknowledges financial support by the B.F. Haley and E.S. Shaw Dissertation Fellowship through a grant to the Stanford Institute for Economic Policy Research.

†Harvard Business School. zcullen@hbs.edu, https://zoecullen.com/.

‡Department of Economics, Stanford University. bph1@stanford.edu, http://web.stanford.edu/∼bph1/.

1 Introduction

Recent policy proposals to mitigate wage inequality include increasing *pay transparency*, the amount of information workers have about each others wages.[1] In many settings, the first step has been to protect the right of co-workers to communicate pay information and extend the time frame during which information from peers is permissible as court evidence of pay discrimination. Following the signing of the 2009 Lilly Ledbetter Fair Pay Act, which removed the statute of limitation for pay discrimination law suits, the federal government has prohibited federal contractors from punishing workers who discuss pay at the workplace, and several states have passed similar laws (including California and Massachusetts in 2016). The stated purpose of protecting the right of workers to discuss pay is ensure that "victims of pay discrimination can effectively challenge unequal pay" through negotiations, by informing them of their employer's willingness to pay for a level of productivity.[2]

However, the debate around pay transparency has paid little attention to equilibrium effects, namely how firms might change their hiring and wage-setting policies in reaction to transparency mandates and how workers might adjust their initial salary negotiations. It also lacks a theory as to why some firms institute transparent pay structures in the absence of any mandate at all. Our research aims to fill this void.

In this paper, we combine empirical analysis of an online labor market, TaskRabbit, that we observe in its entirety over four years, with a parsimonious equilibrium model of dynamic wage negotiations capturing the institutional details of this market. The advent of large online spot markets for labor presents new opportunities for the study of the wage determination process. In an environment largely stripped of career concerns and non-pecuniary benefits, we design direct tests of our model. TaskRabbit matches employers and workers to carry out short-term work in 19 metropolitan areas across the U.S. The institutions in TaskRabbit require workers to bid for tasks but also allow for on-the-job wage renegotiation leading employers to increase worker pay above the initial bids through the

[1] The rise of wage inequality in the United States has been well documented (egs. Katz and Autor, 1999; Saez and Zucman, 2014; Piketty, 2014). Wage inequality has been tied to social costs such as fairness concerns, and far-reaching economic costs, such as inefficient job matching, under-investment in human capital, and even infant mortality (Chen et al., 2014). Citing these economic and social costs, the Equal Pay Act of 1963 and subsequent legislation has increasingly made it illegal for firms to pay workers based on factors other than performance.

These measures have not had the desired effect. In a well-cited statistic on wage inequality, full-time working American women currently earn on average 79% of what their male counterparts make (Blau and Kahn, 2016). The gap is even larger for minority women. This is especially puzzling given that productivity differences between working men and women have largely disappeared. Indeed, the "unexplained" wage gap – the premium that cannot be explained after controlling for observables – is approximately 8%, which is the highest it has been since the wage gap has been noted in modern times (Blau and Kahn, 2016; Goldin, 2014).

[2] https://www.whitehouse.gov/blog/2009/01/25/now-comes-lilly-ledbetter

platform. Pay transparency varies by job setting and by the job posting itself. For example, in some multi-worker jobs, workers are co-located packing boxes in the same office where they might share wage information, and in others they are physically separated distributing marketing materials to different vendors. Some employers choose to use a transparent posted price to advertise their job, and others accept private bids from interested workers.[3]

We assess the costs and benefits of pay transparency along three dimensions: wage equalization, employment rate, and profit sharing between workers and their firm. We find that increasing transparency compresses the wages of similarly productive workers. Some transparency increases employment, but too much reduces employment. Higher transparency shifts expected surplus away from workers and toward the firm. We also find that regulation may be necessary to ensure desirable transparency levels.

To support these claims, we propose an equilibrium model of dynamic wage negotiations between a continuum of workers and a single firm.[4] The level of pay transparency affects the rate at which information about co-worker pay arrives over time to employees. Bargaining takes the form of a directed take-it-or-leave-it (TIOLI) offer from a worker to the firm, as in TaskRabbit, and akin to a phenomenon in the general labor market where workers are asked to put forth their salary expectations to assess whether a match is within the realm of possibilities. Employed workers are able to renegotiate with the firm at will, but workers whose offers are rejected are permanently unmatched with the firm and receive their heterogeneous, exogenous outside options.[5] The firm has a value for labor which is common across workers, but unknown to the workers.[6] Therefore, seeing the pay of a higher paid co-worker is an indication of being underpaid.

We study the unique equilibrium in which the maximum wage offer a firm is willing to accept is a linear function of its value for labor. In it, workers initially bid linear premia over their outside options. Regardless of the level of transparency, workers will only choose to renegotiate their wage once they learn the wage profile within the firm, at which point in time they (successfully) demand pay that is equal to that of the highest earning worker. Therefore, transparency causes an information externality; if a worker finds out that her colleague receives a high wage, she will use this information to negotiate a higher wage for herself. This externality spillover to the way that *initial* wages are set.

There are two major equilibrium effects of increasing transparency: a *demand effect* and

[3]Wage bargaining and transfer of information via transparency are common in many labor markets. Hall and Krueger (2012) find that one-third of workers surveyed explicitly bargain when accepting a job, one-third face posted wages set by their employers, and nearly one-half report that previous wages were used to set current wages. We observe both types of wage-setting, in similar proportions, in TaskRabbit.

[4]We generalize our model to include multiple firms in Appendix D.

[5]We discuss in Section 1.2 how this modeling assumption is derived from our empirical setting.

[6]We discuss in Section 2 how the analysis is unchanged if we instead assume that workers have different productivities but know relative productivity differences.

a *supply effect*. The demand effect is that the firm reduces its willingness to pay for labor when transparency increases because there is a higher chance of an information spillover. In equilibrium with a fully secret pay structure (workers never learn the pay of their peers), the firm accepts all wage offers that are less than the value of labor because there are no information spillovers. With full transparency the firm chooses a posted wage below its value for labor that it pays to all workers, analogously to a monopsonist maximizing its profits. The highest wage the firm is willing to pay for labor is strictly decreasing in transparency.

With increased transparency, the supply effect dictates that workers are willing to accept the job at lower wages initially. The option value of waiting to renegotiate once more information arrives is increasing in the level of transparency, and the premium a worker asks for over her outside option in her initial negotiation is decreasing in the level of transparency and converges to 0.

We first consider the effect of pay transparency on wage equalization. Increasing pay transparency initially (at the time of hiring) increases the wage gap between employed workers with high and low outside options. The supply effect causes all workers to lower their initial wage offers, but low outside option workers reduce initial offers more than workers with high outside options, thus increasing initial wage dispersion. However, as more pay information surfaces over time, each worker becomes more likely to receive the highest wage the firm is willing to offer, equalizing workers' earnings. This latter effect dominates the former in the long run, and so discounted lifetime earnings are compressed as transparency increased.

The combination of demand and supply effects lead to a non-monotonic overall effect of increasing transparency on expected employment level. When transparency is low, the supply effect dominates (fewer workers over-negotiate) which increases expected employment. However, increasing transparency beyond a point means that the demand effect dominates which decreases expected employment. We show that the expected employment level is concave in transparency, meaning that either full secrecy or full transparency is expected employment *minimizing*. An intermediate level of transparency, in which workers learn wage information after joining the firm, maximizes expected employment. This maximizer is falling in both the expected value of labor and the expectation of worker outside options. We also show that a higher level of transparency raises the employment level relatively more when the value of labor is lower.

Pay transparency also changes the division of surplus between workers and firm as it shifts bargaining power away from workers and toward the firm. Under full secrecy workers never renegotiate in equilibrium. Therefore, negotiations involve workers making only an initial offer to the firm. In other words, the workers have the bargaining power and maximize their portion of the surplus, in expectation. On the other hand, under full transparency the

equilibrium outcome is equivalent to the firm selecting the optimal posted wage. Despite the fact that workers are the ones who appear to have the bargaining power due to the construction of the game, full transparency actually shifts this bargaining power to the firm by allowing it to effectively make a TIOLI offer to workers. This allows the firm to maximize its share of expected profits. We find that this intuition holds for intermediate levels of transparency as well. Specifically, we learn that increasing pay transparency increases ex-ante firm profit while decreasing ex-ante worker profits.

A policy maker can select an "optimal" ex-ante level of pay transparency, by weighing these three criteria: pay equality, employment level, and the division of surplus between the firm and workers. However, in reality, the firm exerts some control over information spillovers on site. A firm also has access to more information about its own value of labor than a policy maker. Because of this information asymmetry, it is not inconceivable that firms could endogenously choose a profit maximizing level of transparency in a manner that improves upon the social planner's objectives as well.

We turn our attention to the objective of the firm to maximize profits, and identify the level of transparency selected by the firm in the absence of regulation. Because the profitability of transparency is a function of the firm's value of labor, the firm's choice of transparency signals the firm value to workers, which in turn affects their negotiation tactics, leading to a unique equilibrium outcome in which the firm pools on full transparency regardless of its value of labor. The reasoning for this is unraveling. In any alternative scheme, the lowest value firm type that selects pay secrecy earns zero profits, as workers will always offer no less than this firm type's value for labor in their initial negotiations. This firm type could deviate to full transparency, and post a price below its value to make positive profits, but this would result in a new "lowest value firm type" that receives zero profits when is adheres to equilibrium strategies. This logic unravels toward the firm choosing full transparency for any value.

We use unique back-end data from TaskRabbit from 2010 to 2014 to test our model predictions and study the importance of phenomenon outside our model. We observe all transactions on this platform over this time period, as well as job postings, worker bids, on-the-job bonuses, employer ratings of workers, worker and employer demographics, and cancellations. TaskRabbit staggered its entrance into metropolitan areas, allowing us to analyze the evolution of multiple marketplaces starting from their origin. The dynamics of these labor markets offers a novel window into equilibrium outcomes.

Pay transparency is affected by job and posting characteristics. We identify work settings where physical proximity between multiple workers is unavoidable (high transparency) and those where separation is necessary (low transparency). Employers are also able to endogenously opt into full transparency by posting a price publicly in the job listing.

We find the following in our empirical analysis, consistent with the predictions of our model.

Pay equity: The pattern of pay equalization we observe confirms key assumptions of our model which differ from alternative models of the effect of social concerns on wage compression. Conditional on receiving final pay that is higher than her initial bid for a job, similar workers almost always receive pay equal to that of the highest bidder in a pay transparency environment. At the same time, the distance between workers' bids does not affect the likelihood that a worker receives a raise. Both of these findings are in line with our bargaining story, and we carry out a series of tests that support the causal claim that a communication channel between co-workers about relative pay leads to equilibrium wage equalization.[7] Overall, jobs with partial pay transparency, those with multiple co-located workers, result in final pay that is half as variable as pay in jobs that are otherwise similar, but in which wages are less transparent due to worker separation.

Employment: In the context of TaskRabbit, we measure employment according to whether a job vacancy is filled. Employment is 12% higher among jobs with transparent posted prices.We also show that as employer household income rises (a proxy for increased willingness to pay for labor in this marketplace for chores), the employment gains from higher transparency fall and can become negative.

Profit sharing: While we cannot directly observe the profits of employers in TaskRabbit, we do observe the wage bill. Under transparent pay, holding constant other job factors, the total wage bill is approximately 10% lower. Coupling this with no change in the likelihood of completing a job under higher transparency, we find that higher transparency leads to higher employer profits.

Endogenous, market-level transparency: TaskRabbit staggered entry into metropolitan areas across America and we observe a striking linear progression toward transparent posted wages month-over-month across all markets; for every month on the platform, the fraction of jobs using a transparent posted price in a city increases by 1%. This trend is not explained by the changing composition of jobs or employers on the platform, nor do we find it to be consistent with stories of employer learning. This dynamic unraveling is consistent with the (static) unraveling result from our theory.

[7]We elicit expectations about communication and other activities on the job in a survey of similar online workers. Survey participants report the likelihood that co-workers discuss pay is 47% on average with substantial heterogeneity depending on whether the workers are physically co-located while competing the task.

We seek external validity for our findings. One anomalous institutional feature of TaskRabbit compared to standard labor markets is the simplified bargaining protocol, TIOLI offers, induced by the bidding process on the platform. A concern may be that our findings do not extend to "natural" bargaining settings. To address this concern, we run a randomized field experiment on Amazon's Mechanical Turk. We hire nearly 350 managers to negotiate wages and oversee the efforts of more than 1,000 workers who complete transcriptions of 1940s US census tables. Text-to-text transcription is, a common, real-effort task on Mechanical Turk, allowing us to observe participants in a typical work setting. Each manager is given a budget for each successfully transcribed page, and she keeps the remainder after paying the worker. We vary transparency by restricting wage negotiations to a common chat room containing a manager and multiple workers, or separate workers into private chat rooms with their manager. We do not place restrictions on the bargaining protocol. Our findings are consistent with our theoretical and TaskRabbit findings. Furthermore, by assigning managers' budgets and measuring workers' outside options using an incentive compatible method similar to thatof Becker et al. (1964), we verify theoretical findings with precise parameter estimates.

We combine theory and empirics to run a horse race between our model and a competing theory of social costs of pay inequality. If workers face a morale cost after learning they are underpaid, resulting in low effort, proactive employers may increase the wages of these workers to recuperate high effort.

To assess the impact of morale we build endogenous effort and the morale cost specification of Breza et al. (2016) into our model. Theoretically, only a very extreme and discontinuous morale cost functions could replicate our empirical finding that renegotiation does not depend on the extent of the wage gap, but conditional on renegotiating, wages are equalized. Pay equalization can only be explained within a model including morale costs if workers quit the job (expend 0 effort) upon finding out they are paid even small amounts less than a peer. This is inconsistent with the findings of our field experiment.

Finally, we consider the effects of worker heterogeneity along gender lines. We find that women have outside options that are roughly 10% lower than that of men. We find that the gender pay gap caused by this difference in outside options is mitigated with higher levels of transparency. However, we also find evidence of network effects within jobs: men are more likely to receive bonuses than women when there are communication channels between workers, even within a job, and jobs with more men result in more bonuses overall. One explanation for this is that men are more likely to discuss their wages than women, which we support with survey evidence. We nest these network differences between men and women into our model and show that intermediate levels of transparency can lead to very different de facto arrival rates of wage information for men and women, potentially increasing the

gender way gap. Even with these additions, approaching full transparency levels the playing field on this front and equalizes wages. This evidence may be of interest to proponents of open communication channels about pay within firms as a way to mitigate the gender pay gap.

1.1 Related Literature

At the heart of our paper is the notion that wages may be tied to factors unrelated to productivity. Frank (1984) is an early paper that demonstrates this claim, and there has been a wide literature investigating why this may be. Burdett and Mortensen (1998), Postel-Vinay and Robin (2002), Postel-Vinay and Robin (2006), Cahuc et al. (2006), and Bagger et al. (2014) explain wage dispersion using search models: similar workers naturally receive different job offers over time due to randomness. Those who luck into better offers benefit compared to their unlucky peers. Therefore, wage dispersion in these models is endogenous to the arrival process of offers. In our setting, wage dispersion exists and is endogenous to the bargaining process between worker and employer. Bewley (1999), Abowd et al. (1999), and, more recently, Song et al. (2016) show that wages within firms are compressed relative to the wages across firms due to firm aversion to wage dispersion. In our model, transparency increases the cost of pay inequality within a firm as workers will be able to use this information to negotiate higher wages. We observe significant heterogeneity in the distribution of wages across employers in our data set, which we attribute to exogenous and endogenous differences in bargaining across different employment settings.

Economists have long studied bargaining, and including a comprehensive review of this literature is infeasible. Our paper most closely relates to the double auction literature beginning with a well-known paper by Chatterjee and Samuelson (1983). In this model, a buyer and a seller have private values for a particular good. Both agents simultaneously place bids, and if the bid of the buyer is higher than that of the seller, the two exchange the good at a price that is a predetermined convex combination of the two bids. Satterthwaite and Williams (1989) and Leininger et al. (1989) further examine the equilibria of the double auction game. Radner and Schotter (1989) conduct an experimental study to determine the focality of linear equilibria. The study of simple double auctions stagnated partially because it was determined by Satterthwaite and Williams (1989) that double auctions are generally inefficient means of bargaining within the space of all bargaining mechanisms. Although the double auction model is a static bargaining game, we show a connection to the equilibria of our dynamic game. Therefore, double auctions appear to be a compact representation of a natural bargaining situation, and maximizing desirable properties *within the realm of double auctions* may be relevant to policy creation. Our main results are related to answering this type of question.

There is also a recently active empirical literature on the effects of pay transparency. This literature builds on the fair wage-effort hypothesis introduced to economics by Akerlof and Yellen (1990) which posits a morale cost and associated reduction of effort when a worker believes she is underpaid. Card et al. (2012), Mas (2016), Perez-Truglia (2016), and Breza et al. (2016) conduct field and natural experiments on transparency, and document worker dissatisfaction upon learning of peers with higher pay. Charness and Kuhn (2004), (2007) investigate similar claims in laboratory settings. Nevertheless, these papers do not explicitly investigate how (if at all) workers use information of higher paid co-workers to negotiate higher wages for themselves. We find evidence of this morale effect in a field experiment, but by combining our empirical findings with theory, we show that the observed wage compression is more consistent with a bargaining mechanism than a morale mechanism.

Finally, as pay transparency is often aimed at closing the gender pay gap, our paper fits in to the long literature on gender differences in economic environments. We find a similar gender pay gap in TaskRabbit as in standard economic markets (Blau and Kahn, 2016; Goldin, 2014). We also find network effects between men and women (Keister, 2014) and differences in the likelihood of successful negotiation (Babcock and Laschever, 2003; Niederle and Vesterlund, 2007).

The remainder of the paper is organized in standard fashion. Section 1.2 explains how our modeling decisions fit our empirical environment. Section 2 lays out our theoretical model. Section 3 presents our main theoretical findings. Section 4 exhibits empirical tests of our main findings using TaskRabbit data. Section 5 discusses our field experiment and related findings. Section 6 examines an alternative model based on the fair wage-effort hypothesis and morale costs associated with pay transparency. Section 7 investigates heterogeneity of workers across gender lines, and the effects of transparency on the gender pay gap. Section 8 concludes. Omitted proofs and regression tables are contained in the Appendix.

1.2 Links Between Empirical Setting and Model

The initial wage agreement between a worker and employer provides a binding floor, which can be increased through re-negotiation at any time during the course of employment. Matching this feature, in our model workers place take-it-or-leave-it offers at the onset of the negotiation.

Most jobs on TaskRabbit are chores, and as such, labor is relatively homogeneous and low-skill. We model this aspect by endowing all workers in our model with the same (unknown) level of productivity. Our results do not change if we instead assume that workers have different productivity levels, but that their relative productivities are observable.

In the context of a richer labor market, additional features will also affect relative bargaining power, such as the extent that search frictions prevent a worker from applying for

an equivalent job, or the employer from hiring a substitute worker without a delay. Jobs on TaskRabbit generally have imminent deadlines and fill quickly; 97% of tasks are completed within three days of posting. Therefore, if an on-the-spot negotiation fails, search frictions are high. As a consequence, the ability for either side to find an immediate replacement is low. Furthermore, these jobs are temporary, irregular tasks, meaning that most workers and employers have jobs off platform; Farrell and Greig (2016) find that online labor platform users earn on average one-third of their total income on platform. Cullen and Farronato (2016) show workers supply labor flexibly in response to demand fluctuations and earn close to their cost of time. To reflect these traits in our model, a worker's only alternative to the job is a fixed outside option and the firm's only alternative is an outside option of 0 for each worker it fails to employ.

With these modeling decisions explained by our empirical setting, we proceed to formally define the game.

2 Model

2.1 Preliminaries

Our model, and specifically the bargaining portion of it, is based on TaskRabbit. Time is continuous, and is indexed by $t \in \mathbb{R}_+$. There is a single firm in the economy.[8] At each time t, a ρ mass of workers enters the market, and existing workers exogenously depart the market via a Poisson process with rate ρ. Each worker i has an (private) outside option $\theta_i \sim G[0,1]$ i.i.d., which is the flow payment i receives when not matched to the firm. Productivity of labor is $v \sim F[0,1]$, which is common across workers, and is known only to the firm.[9] The firm faces a constant returns to scale production function, and receives a flow surplus v from each employed worker. All agents discount the future at rate δ, are risk neutral, and seek to maximize discounted expected flow payments. We assume that F and G are C^2 over the interior of their supports.

At $t = 0$ (and before any workers enter the market) the firm selects a maximum wage it is willing to pay for a worker as a fuc$\bar{w}(v) \in [0,1]$, where the choice of $\bar{w}$ is not immediately observed by workers. Each worker bargains for wages by making take-it-or-leave-it (TIOLI) offers to the firm at any point during her employment, potentially renegotiating infinitely often. Workers who at anytime offer a wage greater than $\bar{w}$ are permanently unmatched with

[8] We generalize our model to include multiple firms in Appendix D.

[9] The assumption of a common value of labor is made for ease of exposition, and the analysis relies only on the assumption that workers can observe their relative productivity within the firm. We could complicate the analysis by supposing each worker has a known type $\tau \in \mathcal{T}$ where $\mathcal{T}$ is some countable set. Let $v_\tau \sim F[0,1]$ i.i.d. but unknown to workers. The analysis is almost entirely unchanged, other than additional notation, with this modification.

the firm. Let W_t denote the set of wages the firm is paying to its employed workers, where $W_0 = \{\bar{w}\}$. We model transparency as a random arrival process; at time t matched workers observe W_t according to an independent Poisson arrival process with rate $\lambda \in [0, \infty) \cup \{\infty\}$, where we take $\lambda = \infty$ to mean that the process arrives at every time t. For much of the paper, we abstract away from the genesis of this arrival process. It can be thought of as the frequency with which there is an information leak, that existing workers see the offers of incoming workers, or wage gossip between workers.[10]

The timing of the stage game is as follows at each time $t \geq 0$:

1. New workers enter the market and are matched with the firm.
2. Each matched worker i learns W_t independently with arrival rate λ.
3. Newly entering workers must bargain with the firm and any existing, matched worker can initiate bargaining. Any worker i who engages in bargaining makes a TIOLI offer $w_{i,t} \in [0, 1]$ to the firm. If $w_{i,t} \leq \bar{w}$ the firm pays i a flow wage $w_{i,t}$ until i departs or attempts to renegotiate. If $w_{i,t} > \bar{w}$ then the firm unmatches with i and the two are ineligible to rematch.
4. Existing workers depart at rate ρ.

In Appendix B.1 we discuss a generalization of the game in which the firm accepts or rejects offers as they arrive, and show that all of the results of this paper are unchanged (under a proper selection of "Markov" equilibria.) In Appendix D we expand our model to allow workers to search for work between multiple firms, and show that many results are robust to this extension. In Appendix E we discuss alternative bargaining protocols and show that our results are robust to these settings.

2.2 Equilibrium Selection

We investigate pure strategy perfect Bayesian Equilibria of the game. We restrict our attention to equilibria satisfying the following conditions as a method to simplify the exposition:

A1 $0 \leq \bar{w} \leq v$. Let w_i^* be the initial wage offer of worker i during the first period she meets the firm. If $v \leq w_i^*$ for every worker i then $\bar{w} = v$.

A2 $\theta_i \leq w_i^* \leq 1$. If there is no v such that $\theta_i \leq \bar{w}$ then $w_i^* = \theta_i$.

A3 $\bar{w}$ and w_i^* are strictly increasing and continuous functions of v and θ_i, respectively

[10] We discuss endogenous information arrival in Section 7.

A4 Off path, each employed worker i believes with probability 1 that $\bar{w} = \theta_i$ until i receives wage information.

A5 If worker i initiates re-negotiation at time $t' > t$ then $w_{i,t'} > w_{i,t}$.

A1 and **A2** restrict actions for agents who never match on equilibrium path. Low value firms who can never afford to profitably hire any workers will set $\bar{w} = v$. Similarly, workers with high outside options who know any initial offer greater than their outside options will not be accepted in equilibrium will ask for exactly their outside options. These assumptions rule out pathological equilibria, such as one in which the firm sets $\bar{w} = 0$ and every worker offers $w_i^* = 1$.

A3 limits our study to continuous equilibria. This assumption is made for tractability and removes "step function" equilibria, in which workers and the firm pool on a predetermined wage, from consideration. Leininger et al. (1989) suggest similarities between the set of continuous equilibria and the set of step function equilibria of a game similar to our own, and so we do not believe this to be a conceptually limiting constraint.[11]

A4 pins down off-path beliefs. It is easy to see that these are the most favorable beliefs to the firm allowable in a PBE. In the event that an unanticipated offer is accepted, the worker believes she is extremely lucky and is receiving the highest possible wage until she is presented with evidence to the contrary. Therefore, any equilibrium sustainable under these off-path beliefs is sustainable under any set of off-path beliefs.

A5 rules out a multiplicity of essentially equivalent equilibria by restricting that a worker who attempts to renegotiate must demand a strictly higher wage that she was previously receiving. Otherwise, for example, a worker could "renegotiate" infinitely often but just offer the same wage over and over again.

2.3 How do workers renegotiate?

Workers learn the wages of their co-workers over time and are able to initiate bargaining infinitely often if they wish. There is reason for concern in dynamic relationships of a "ratchet effect" (Weitzman, 1980) – that successful negotiation will lead to an increased desire on the part of workers to initiate bargaining in future periods. We find, however, that in equilibrium, a worker will only negotiate at most twice during her tenure at the firm, upon first being matched to the firm (when she will make a wage offer dependent on her outside option), and upon learning the highest wage paid within the firm. Upon re-negotiating, the worker will receive $\bar{w}$ for the duration of her tenure at the firm.

[11] We discuss this similarity in Section 2.4.

PROPOSITION **1.** *In equilibrium workers only negotiate wages in the first instant they are hired and the first instant they receive information about wages of co-workers. Upon renegotiating, workers offer and receive* $\bar{w}$.

The intuition for this result is that workers do not learn about v (and hence do not learn about $\bar{w}$) if their offer is accepted. Any worker strategy that says "offer w when initially hired at time t and offer $w' > w$ at time $t' > t$ if I have not learned the wages of my co-workers" cannot be optimal, because even if offering w' at time t' improves the expected utility of the worker, she would have been better off asking for w' at time t as $\bar{w}$ is fixed over time. This reasoning is shared in Tirole (2016), albeit in a different context.

Due to the continuum of workers entering the market at each time t and assumptions **A1-3**, workers trace out the entire set $[0, 1]$ with their initial offers at each time t. Therefore, the highest wage paid by the firm is $\bar{w}$ for all $t \geq 0$. As a result, the maximum wage that a worker observes upon information arrival is $\bar{w}$. Clearly, a worker will then demand this amount from the firm.

2.4 Equilibrium conditions

Given that workers negotiate at most twice in equilibrium, we are able to use standard techniques to solve for the optimal bargaining strategies of agents. Letting $\bar{F}(x) = P(\bar{w} \leq x)$, worker i negotiates at the first moment she is hired to:

$$\underset{w_i^*}{\operatorname{argmax}} \left(\frac{w_i^*}{\rho + \delta + \lambda} + \frac{\lambda}{\rho + \delta + \lambda} \frac{\mathbb{E}\left(\bar{w} | \bar{w} \geq w_i^*\right)}{\delta + \rho} \right) \left(1 - \bar{F}(w_i^*)\right) + \frac{\theta_i}{\delta + \rho} \bar{F}(w_i^*) \tag{1}$$

where the first term represents the weighted (by λ) expected wage the worker receives if matched with the firm, and the second term represents the lifetime earnings of the worker if she exceeds $\bar{w}$ and instead consumes her outside option for her lifetime. When $\lambda = \infty$, the pricing scheme is a posted price in which all workers can elect to make an offer $w_i^* = \bar{w}$ or unmatch with the firm.

In a series of steps (shown in Appendix B.2), we modify the objective function without affecting the maximizer, and show that this is equivalent to solving:

$$\underset{w_i^*}{\operatorname{argmax}} \int_{w_i^*}^{1} \left((1 - \Lambda) w_i^* + \Lambda x - \theta_i \right) \bar{f}(x) dx \tag{2}$$

where $\Lambda = \frac{\lambda}{\rho + \delta + \lambda}$ for all $\lambda \in [0, \infty)$ and $\Lambda \equiv 1$ for $\lambda = \infty$. In equilibrium, the firm sets $\bar{w}(v)$ to

$$\underset{\bar{w}}{\text{argmax}} \frac{\int_0^{\bar{w}} (v-y)\,\bar{g}(y)dy}{\rho+\delta+\lambda} + \bar{G}(\bar{w})\frac{\lambda}{\rho+\delta+\lambda}\frac{1}{\rho+\delta}(v-\bar{w}) \tag{3}$$

where $\bar{G}(x) = P(w_i^* \leq x)$. The first term gives the total discounted profits made by the firm before the workers experience an event (seeing the wage profile or perishing) and the second term is the profit made from workers after renegotiating their wages to $\bar{w}$ over the rest of their lifetimes in the firm. We can similarly manipulate the objective as with the worker problem:

$$\underset{\bar{w}}{\text{argmax}} \int_0^{\bar{w}} (v-(1-\Lambda)\,y-\Lambda\bar{w})\,\bar{g}(y)dy \tag{4}$$

These manipulations map our problem into the well-known double auction of Chatterjee and Samuelson (1983) in which a seller (worker) with a private value for a good (θ_i) and a buyer (firm) with a private value for a good (v) submit sealed bids. If the bid of the buyer is at least as large as that of the seller, the good switches hands at a price set be a predetermined convex combination of the two bids (determined by Λ). The first order conditions for the worker and firm are

$$w_i^* - \theta_i = (1-\Lambda)\,\frac{1-\bar{F}(w_i^*)}{\bar{f}(w_i^*)} \tag{5}$$

and

$$v - \bar{w} = \Lambda\frac{\bar{G}(\bar{w})}{\bar{g}(\bar{w})}. \tag{6}$$

We further know from Satterthwaite and Williams (1989) that the set of equilibria exactly correspond with solutions of the first order equations for distributions F and G with strictly increasing virtual values, that is, $\theta + \frac{G(\theta)}{g(\theta)}$ is strictly increasing in θ and $v - \frac{1-F(v)}{f(v)}$ is strictly increasing in v.

2.5 Solving for equilibrium

The optimal bidding and wage setting policies of the firms and workers are interdependent; workers decide how aggressively to bid depending on how the firm sets $\bar{w}$, while the firm sets $\bar{w}$ as a function of how aggressively the workers bid. Satterthwaite and Williams (1989) show that there exists a continuum of equilibria satisfying Equations 5 and 6. Our set lacks natural ordering, limiting the possibility for general claims about the entire set of equilibria, but we build from experimental evidence suggesting that equilibria in which w_i^* and $\bar{w}$ are linear functions of θ_i and v, are focal Radner and Schotter (1989). We produce similar evidence

in a setting similar to TaskRabbit in Figure 9. Therefore, we focus our analysis on linear equilibria, and restrict attention to a two parameter family of power law distributions of worker outside options and firm values, which we show admit a unique linear equilibrium.[12] We then study the properties of this equilibrium, and analyze the effects of transparency.

The class of distributions we study are:

$$\begin{aligned} F(v) &= 1-(1-v)^r, \quad r>0 \\ G(\theta) &= \theta^s, \qquad\qquad\quad s>0 \end{aligned} \tag{7}$$

As r increases, v is on average lower and as s increases, θ is on average higher. Therefore, increasing r or s reduces the average benefits from employment. We formally define a linear equilibrium below and then demonstrate that distributions within this family admit a unique linear equilibrium.

DEFINITION **1.** *A* linear equilibrium *is a pure strategy perfect Bayesian equilibrium satisfying* ***A1-5****, where $\bar{w}$ is a linear function of v whenever a positive mass of workers offers $w_i^* \leq v$, and where w_i^* is a linear function of θ_i whenever there is positive probability that $\theta_i \leq \bar{w}$.*

PROPOSITION **2.** *For any pair of distributions within the family described in Equation 7 there exists a unique linear equilibrium.*

What can we say about the equilibrium bargaining strategies of workers and firms, and how are they affected by transparency? First, equilibrium wages will fall in some interval $[a,h] \subset [0,1]$. This means that firm types with values below a will not hire any workers, and all workers with outside options above h will remain unemployed. Second, we can see from Equations 5 and 6 that all workers and firm types with positive probability of matching in equilibrium charge a premium, that is, $v - \bar{w} \geq 0$ and $w_i^* - \theta_i \geq 0$. We show that higher outside option workers and lower value firm types face higher risks of being unmatched in equilibrium. As a result, the markup charged by each worker is decreasing in θ_i and the markdown set by the firm is increasing in v. We further can show that both $\bar{w}$ and w_i^* are decreasing in Λ; with increased transparency the firm reduces the highest worker offer it accepts to avoid information spillovers across workers (which we call the **demand effect**), and workers make more conservative initial offers because with higher transparency the option value of waiting to receive a high wage is increasing relative to the risk of securing a high initial wage (which we call the **supply effect**). The following proposition formalizes these arguments.

[12] The approach of making parametric assumptions to ensure linear equilibrium is common. One recent example on CEO pay is Edmans et al. (2016). Power law distributions are commonly observed in economic situations such as ours, including worker income and firm productivities. See Gabaix (2009, 2016) for details.

PROPOSITION **3.**

1. $v - \bar{w} \geq 0$ *and strictly increasing in* v *for all* $v \in [a, 1]$,

2. $w_i^* - \theta_i \geq 0$ *and strictly decreasing in* θ_i *for all* $\theta_i \in [0, h]$,

3. $\bar{w}$ *is strictly decreasing in* Λ *for all* $v \in [a, 1]$. *As* $\Lambda \to 0$, $\bar{w} \to v$ *for all* $v \in [0, 1]$, *and*

4. w_i^* *is strictly decreasing in* Λ *for all* $\theta_i \in [0, h]$. *As* $\Lambda \to 1$, $w_i^* \to \theta_i$ *for all* $\theta_i \in [0, 1]$.

The decline of $\bar{w}$ in Λ is similar to the strategy of a monopsonist that optimally limits demand. Due to information externalities, the firm reduces $\bar{w}$, thus restricting the extensive margin of labor (the number of workers it hires) while increasing the intensive margin (profit per worker hired). There are two special cases that further clarify this point. Consider full secrecy ($\Lambda = 0$) and full transparency ($\Lambda = 1$). In the full secrecy case, there are no information spillovers. Therefore, in the unique equilibrium the firm must set $\bar{w} = v$, meaning that the firm sets the highest acceptable wage equal to the marginal product of labor. In the full transparency case, there are perfect information spillovers, and every worker learns the wages of others within the firm at the instant they are hired. Since all workers will learn $\bar{w}$ before their first negotiation, this is equivalent to posting a profit maximizing wage $\bar{w}$ given a supply curve of labor (the distribution of outside options), which is the exact problem that a traditional monopsonist faces. If the firm instead sets $\bar{w} = v$ as in the full secrecy case, it would earn zero profits.

We graphically represent the **demand** and **supply effects** in Figure 3 as Λ increases from 0 to $\frac{3}{4}$.

3 Main results - Effects of transparency on equilibrium

We analyze the equilibrium effects of transparency along three dimensions: Does transparency lead to pay equity? Does transparency increase employment? and Does transparency benefit workers or firms?

3.1 Income inequality

Our theoretical framework allows employees and workers to change their behavior in anticipation of future communication about pay. Indeed, given the renegotiation that occurs once new information about co-worker pay is leaked, workers will initially offer lower wages for the task at hand to reap higher benefits later from the renegotiation.

In our theoretical analysis of wage equalization we compare only workers i and j who are hired under both Λ and $\Lambda^{'}$ so we do not confound compression effects with selection effects.[13] For any two workers i and j with $\theta_i > \theta_j$ who are hired under two levels of transparency $\Lambda < \Lambda^{'}$, there are two effects. First, as shown in Proposition 3, the **supply effect** incentivizes agents to reduce initial wage offers. We show that in equilibrium, since j has a lower outside option than i, j reduces her initial offer more than i. For intuition, we see in Figure 3 that the relative impact of the supply effect on w_i^* is smaller the larger θ_i is. Second, increasing transparency augments the rate at which both workers receive $\bar{w}$, reducing the dispersion in the lifetime earnings of workers as an increase in salary to $\bar{w}$ is a larger raise the smaller θ_j is. The first effect increases the initial wage gap between i and j in the short run, however, we show that the latter effect dominates in the long run, leading to more compressed expected lifetime earnings with higher transparency. We document these two effects by plotting the expected difference in wages between workers i and j over time and for different levels of transparency in Figure 4.

THEOREM **1.** *Let $\theta_i > \theta_j$, $1 > \Lambda^{'} > \Lambda^{''}$, and suppose workers i and j are both hired in equilibrium under $\Lambda^{'}$ and $\Lambda^{''}$.*

1. *The difference in initial offers $w_i^* - w_j^*$ is higher under $\Lambda^{'}$ than $\Lambda^{''}$ and*
2. *Let $T(\Lambda, v, \theta_k)$ be the equilibrium expected discounted lifetime earnings of a worker k with outside option θ_k under transparency level Λ and firm value v conditional on k being employed at the firm. Then $T(\Lambda^{'}, v, \theta_i) - T(\Lambda^{'}, v, \theta_j) < T(\Lambda^{''}, v, \theta_i) - T(\Lambda^{''}, v, \theta_j)$ and $T(\Lambda^{'}, v, \theta_i) - T(\Lambda^{'}, v, \theta_j) \to 0$ as $\Lambda^{'} \to 1$.*

Note that $w_i^* - w_j^* = 0$ and $T(\Lambda^{'}, v, \theta_i) - T(\Lambda^{'}, v, \theta_j)$ and for any two employed workers i and j when $\Lambda^{'} = 1$, as all workers earn $\bar{w}$ in the moment they are hired. While increasing transparency (within the interior) increases initial wage dispersion, it has the effect of compressing discounted lifetime earnings. The long-run effect of transparency, therefore, is to equalize wages of workers with different outside options. The latter point is easiest to understand in the context of a Samuelson Chatterjee double auction. Increasing Λ reduces the

[13] Does increasing transparency from Λ to $\Lambda^{'}$ compress the expected equilibrium lifetime earning of workers *employed within the firm under both Λ and $\Lambda^{'}$*? The caveat is necessary as we show in Theorem 2; increasing transparency can increase employment, meaning that a previously unemployed, high outside option worker may find employment only when transparency is increased.

To make this point concrete, take some small $\epsilon > 0$ and consider increasing transparency from Λ to $\Lambda^{'} = \Lambda + \epsilon$, such that more workers are employed in equilibrium under $\Lambda^{'}$. In Appendix C we show that w_i^* and $\bar{w}$ are continuous in Λ and so the expected lifetime earnings of any worker j hired under both transparency regimes is barely affected by an ϵ increase in transparency. However, a worker i who over-negotiates at level Λ receives her outside option θ_i for her entire lifetime, while if she manages to find employment at the firm under $\Lambda^{'}$ her average lifetime earnings will be greater than and bounded away from θ_i (as she always asks for a premium $w_i^* - \theta_i > 0$). But note that $\theta_i > \theta_j$, so the lifetime earnings of i and j are not compressed by increased transparency.

bargaining power of the workers, and therefore, the different bargaining positions of workers with heterogeneous outside options make up a smaller proportion of their expected lifetime earnings.

The previous theorem deals with wage compression. This may be a reasonable notion of fairness when θ_i represents a worker's outside option. However, compression of expected surplus, that is, the difference between wage and θ_i, is another notion of fairness which may be particularly relevant to consider in cases when θ_i represents the flow cost a worker bears for completing the job.[14] Does transparency lead to a compression of expected surplus across workers? Simply put, the wage compression results in Theorem 1 are reversed when we consider surplus compression as low θ_i workers are those who enjoy the largest surplus under employment. We state these results formally below.

COROLLARY **1.** *Let $\theta_i > \theta_j$, $1 > \Lambda' > \Lambda''$, and suppose workers i and j are both hired in equilibrium under Λ' and Λ''.*

1. *The difference in initial surplus $(w_j^* - \theta_j) - (w_i^* - \theta_i)$ is smaller under Λ' than Λ'' and*
2. *Let $S(\Lambda, v, \theta_k)$ be the equilibrium expected discounted lifetime surplus of a worker k with outside option θ_k under transparency level Λ and firm value v conditional on k being employed at the firm. Then $S(\Lambda', v, \theta_j) - S(\Lambda', v, \theta_i) > T(\Lambda'', v, \theta_j) - T(\Lambda'', v, \theta_i)$, and $T(\Lambda', v, \theta_j) - T(\Lambda', v, \theta_i) \to \frac{\theta_i - \theta_j}{\rho + \delta}$ as $\Lambda' \to 1$.*

3.2 Employment

Consider an increase in transparency from Λ to Λ'. In an abuse of notation, let $\bar{w}_\Lambda$ denote the maximum willingness of the firm to pay and $w_{i,\Lambda}^*$ the initial offer of worker i for transparency level Λ. The **demand effect** lowers employment as $\bar{w}_{\Lambda'} \leq \bar{w}_\Lambda$ means that there are fewer workers with $\theta_i \leq \bar{w}_{\Lambda'}$ who are eligible for employment. The **supply effect** increases employment as $w_{i,\Lambda'}^* \leq w_{i,\Lambda}^*$ for all i meaning that fewer workers who over-negotiate by initially offering $w_{i,\Lambda'}^* > \bar{w}_\Lambda$. As such, the effect of increasing Λ on employment is neither obvious nor (necessarily) monotonic.

THEOREM **2.** *Ex-ante equilibrium employment is concave in Λ and maximized at*

$$\Lambda^* = \frac{1 - \mathbb{E}(\theta)}{1 + \mathbb{E}(v) - \mathbb{E}(\theta)} \tag{8}$$

and the ex-post employment maximizing level of transparency is weakly decreasing in v.

[14] We show in Section 6 that the equilibrium outcome in a game in which all workers have an outside options equal to zero and θ_i represents a flow cost to being employed is identical to that of the game presented above.

The expected surplus generated in a match between a worker and firm is $\mathbb{E}(v) - \mathbb{E}(\theta)$, so Equation 8 tells us that the ex-ante employment optimizing level of transparency is

$$\frac{1 - \text{expected outside option}}{1 + \text{expected match surplus}}$$

We see several important effects of transparency on employment. First, an interior level of transparency (neither full secrecy nor full transparency) maximizes employment. In fact, due to the concavity of the employment in Λ either full secrecy or full transparency is employment *minimizing*.

Second, Λ^* is decreasing in both $\mathbb{E}(v)$ and $\mathbb{E}(\theta)$. Indeed, as $\mathbb{E}(v)$ converges to 0 full transparency becomes close to employment maximizing, and as $\mathbb{E}(\theta)$ converges to 1 full secrecy becomes close to employment maximizing. For intuition, we return to Proposition 3. As $\mathbb{E}(v)$ decreases, the firm's markdown $v - \bar{w}$ is likely to be small regardless of Λ. Therefore, increasing transparency does not greatly reduce the number of workers with $\theta_i < \bar{w}$. But by increasing transparency, workers will shade down their initial offers w_i^*, reducing the number of workers who over-negotiate. Similarly, as $\mathbb{E}(\theta)$ increases, most workers are offer small premia $(w_i^* - \theta_i)$ regardless of Λ. The effect of increasing transparency does little to affect this premium, but instead discourages the firm from setting a large markdown.

Third, both the ex-ante and ex-post optimal levels of transparency are in some sense decreasing in v. We have already discussed how Λ^* is strictly decreasing in $\mathbb{E}(v)$, and increasing transparency is more beneficial for the employment level when v is small. Given that workers' initial offers are not affected by the realization of v (they do not observe it), high transparency causes firms to reduce $\bar{w}$ more significantly when v is high, leading to relatively less employment. Additionally, this comparative static on the ex-post employment maximizing level of transparency also holds for the ex-post *social surplus* maximizing level of transparency. In fact, the ex-post maximizer of employment also maximizes ex-post social surplus; because each employed worker earns a wage weakly greater than her outside option, in equilibrium each employed worker increases social surplus by $v - \theta_i > 0$, implying that social surplus is proportional to employment level. Therefore, increasing transparency is also more beneficial from a social surplus perspective when v is small.

3.3 Profit share: who benefits from transparency?

We now turn our attention to the profit share. Does increasing pay transparency increase worker or firm utility? Indeed, in light of Theorem 2 we may suspect that the efficiency gains caused by transparency could make both firms and workers better off. Nevertheless, we find, perhaps counter intuitively, that increasing pay transparency increases the expected profit of the firm while decreasing the expected welfare of workers. Note that these statements are

all made before the realization of v or θ_i for an agent.

THEOREM **3.** *The ex-ante expected equilibrium profit of the firm is strictly increasing in* Λ *and the ex-ante expected equilibrium profit of workers is strictly decreasing in* Λ.

At first glance this result may seem incorrect. After all, increasing Λ increases the rate at which workers receive the wage $\bar{w}$. However, by increasing transparency, workers lower w_i^* and the firm lowers $\bar{w}$, the effect of which is to lower wages. The overall effect benefits the firm at the expense of workers. For clear intuition, consider the extreme cases of full secrecy and full transparency. Under full secrecy workers have all the bargaining power, as each worker makes a once-and-for-all offer to the firm. Under full transparency, the firm selects a maximum wage which all workers earn in equilibrium the moment they are first hired, which results in the firm having all the bargaining power. Myerson (1981) and Williams (1987) show that the equilibrium outcome under $\Lambda = 1$ maximizes firm ex-ante expected profits and minimizes worker expected surplus within the class of bargaining mechanisms that respect the incentive constraints of workers (the wage cannot fall below θ_i for each i) and that the equilibrium outcome under $\Lambda = 1$ maximizes ex-ante expected worker surplus and minimizes firm expected profit of any bargaining mechanism that respects the incentive constraint of the firm (the salary of a worker cannot be greater than v). We show that this intuition holds as we increase Λ within $(0, \infty)$.

We do not view the shift of profit from workers to firm as an inherently good or bad thing.[15] However, there may be macro-level effects of changing the profit split that we do not capture in our model. We point out this effect of pay transparency and leave its consequences on other economic measures to future work.

3.4 Endogenous firm transparency choice

Until now, we have been studying the effects of increasing transparency ex-ante (before seeing the draw of v) on employment, wage compression, and profit share. These results give insights into the effects of governmental policy instituting pay transparency measures. Although there has been recent state and federal action in the United States to increase pay transparency, many industries remain unregulated. Without regulations, the firm can select transparency to maximize its profits after seeing the draw of v. Furthermore, workers may not directly observe the level of transparency selected by the firm.[16]

[15] Indeed, if we adopt the point-of-view that workers own the firm (perhaps in proportions that are not correlated with outside options) then the split of profit is likely not a first-order concern.

[16] It is likely to be viewed as cheap talk for an employer to say, "our firm has a high degree of transparency, so don't worry, you're likely to learn the wages of your coworkers very soon" at a job interview. Furthermore, workers may not even have all information on firm policies at the time that they accept a job. For example, Marx (2011) finds that firms strategically wait until after workers have negotiated initial salaries and begun working to introduce non-compete agreements.

We show that in contrast to Theorem 3, full transparency is not the profit maximizing (exogenous) level of transparency for every draw of v.

EXAMPLE **1.** *Let $v = 1$ and let $\mathbb{E}(\theta) = \mathbb{E}(v) = \frac{1}{2}$. This implies that $r = s = 1$. We can calculate the profit $\pi(v, \Lambda)$ of the firm using Equation 27 in the Appendix. We see that $\pi(1, 1) = \frac{1}{2}$ while $\pi(1, \frac{1}{2}) = \frac{9}{16}$. Therefore, full transparency does not necessarily maximize the profit of the firm for every v.*

Indeed, the profit maximizing level of Λ is not even monotonic in v. Nevertheless, when firms are able to endogenously select transparency, we find that the unique equilibrium that does not involve employed workers renegotiate in the absence of learning the wages of their coworkers is one in which the firm selects full transparency *regardless of its draw of* v. As wage negotiations between any worker and her employer are relatively rare,[17] we believe this to be a reasonable class of equilibria to study. Our findings suggest that in the absence of governmental regulation, observed levels of transparency may be very different than employment maximizing levels.

THEOREM **4.** *When the firm can privately select Λ as a function of v, there is an essentially unique equilibrium outcome in which each no worker renegotiates in the absence of coworker wage information on equilibrium path. In equilibrium, the firm selects $\Lambda = 1$ for all $v > 0$.*

Because workers do not necessarily observe the transparency choices of firms, each worker's belief of the selected Λ decreases continuously in the length of time since being hired that the worker does not learn the wages of coworkers. When this happens, workers will renegotiate their wages in the absence of learning the wages of their coworkers, reflecting the new information they gain over time. This rules out any strategies in which the firm selects any $\Lambda \in (0, 1)$. Can it be the firm selects $\Lambda \in \{0, 1\}$ in equilibrium? We show that the firm can never set $\Lambda = 0$ in equilibrium due to unraveling. To see this, let v_L be the infemum value for which the firm selects $\Lambda = 0$. Then upon arriving at the firm, the workers will immediately deduce that the firm has selected $\Lambda = 0$ since they do not initially observe the wage profile of the firm. Workers will infer that the firm's value is at least v_L, and so every worker will bid at least v_L. As a result, when the value of the firm is (close to) v_L it will make (approximately) 0 profits unless it deviates to selecting $\Lambda = 1$. But once the firm deviates, there is a new "v_L." Following this logic inductively shows that any plan in which there is a positive measure of firm types playing $\Lambda = 0$ cannot be an equilibrium. The equilibrium in which the firm selects $\Lambda = 1$ for all v can be supported with the off path beliefs that a deviating firm has value $v = 1$ with probability 1. Since $\bar{w} = v$ when $\Lambda = 0$, a deviating firm will make zero profits.

[17] Hall and Krueger (2012) find that about 70% of workers have not negotiated raises at their current jobs.

This result is particularly applicable to online labor markets in which the level of transparency employers can select is very coarse. For example, in our data sample from TaskRabbit, employers of single-worker tasks have the option of accepting bids or posting a price. As these are single-worker jobs, there is no transparency (there is no chance of talking to co-workers) and so accepting bids is equivalent to setting $\Lambda = 0$. As we have discussed earlier, posting a price is akin to full transparency, i.e. $\Lambda = 1$. In such a world, the unraveling result holds without any caveats.

COROLLARY **2.** *When the firm can select $\Lambda \in \{0, 1\}$ as a function of v there is an essentially unique equilibrium outcome. In equilibrium, the firm selects $\Lambda = 1$ for all $v > 0$.*

4 Study I: Evidence from TaskRabbit

4.1 Platform

We use administrative data from an online labor platform, TaskRabbit, between June of 2010 and May of 2014. TaskRabbit differentiates itself from other online labor platforms by specializing in local jobs, which account for 89% of jobs completed. The platform was active in 19 U.S. metropolitan areas across the U.S. during this period. To participate in the marketplace, workers must pass a criminal background check and cursory screening to join the platform and employers must enter a valid credit card.

Our research concentrates on jobs that are posted as one-time tasks.[18] Approximately half of all employers are incorporated businesses and the remainder are designated as households. The employer can observe workers' profiles, which include the number of prior jobs completed on the platform, a rating out of five stars and a short bio.

Employers post a description of the task, details about the exact location, a preferred start time, and a deadline for completion. Workers search through these postings and submit bids to complete the entire project. Alternatively, the employer can choose to post a take-it-or-leave it price, and the first worker to accept is matched.

4.2 Bargaining Environment

Our model of bargaining is premised on TaskRabbit. Workers submit private bids to the employer, and the employer can either accept or reject those bids. On TaskRabbit the employer also has an additional option available, to post a take-it-or-leave it price that all

[18]We classify one-off tasks two ways. In the main specification we limit tasks to those the employer indicates as "non-repeating" when filling out the vacancy forms, and as a robustness check we include the survey responses of several thousand Mechanical Turk workers who read the job descriptions and answered the question, "what is the likelihood that a worker could be rehired for a similar job by this employer?"

workers can observe. By default, the agreed upon wage is released to the worker when she declares the job complete. Alternatively, and regardless of the way in which initial pay is set, employers can elect to increase the wage of any worker through the platform once the job is completed. This allows for the possibility of on-the-job bargaining between worker and employer for a higher wage.

Once the job begins, several frictions make canceling costly for both parties. TaskRabbit is a spot market designed for urgent tasks; half of all posted jobs include a deadline within three days from the time the job is posted, and only 10% of jobs can be postponed beyond two weeks. On average, the median job receives 1 offer in the first hour per vacancy, and 1 every 4 hours over the first day. Taken altogether, we deduce that finding a replacement worker once the job has begun would frequently result in costly delays.

Similarly, workers cannot costlessly transition to another job. Because these are in-person tasks, we observe high travel costs relative to the final transaction price.[19] At the time that a worker is assigned a job, the worker and employer enter a formal contract that can only be cancelled by contacting the platform and providing a reason. TaskRabbit has a three strike rule. After three cancellations a user will not be permitted to use the site again. About 8% of assignments are canceled. During the window between when the match is made and the job is complete, money is held in escrow.[20]

4.3 Measuring Transparency

We measure pay transparency on TaskRabbit three ways. Our first measure is whether the job post itself includes a take-it-or-leave-it posted price publicly visible to all workers. The posted price can either be text embedded in a job description or a public price associated with the job posting format selected. We classify these as full transparency.

Our second measure is based on the physical proximity of workers in multi-worker jobs and the length of time they overlap in the same location. We distinguish settings inherently suited for either co-located workers and physically separated workers, for example, a retail branch might outsource the boxing of holiday gifts at the store (co-located workers) and or outsource the distribution of catalogues in different neighborhoods (separated workers). We use GPS coordinates to classify proximity and we supplement it with survey evidence. We hire approximately five thousand online workers to read through the detailed job descriptions

[19] We calculate the distance between worker residence of work location.

[20] The platform reserves the right to revoke user privileges should any activity suggest circumventing the online contract, including suggestive comments over the platform. We do not rule out the possibility that the relationship between the employer and employee can continue off the platform, we do assume that the matches consecrated over the platform results in a transaction over the platform. For robustness, we replicate results to exclude and include employers that never return to the site for confirmation. Between 47% and 52% of users from each work category do not return to the platform at the time that the job concludes and confirmation is requested.

and report key attributes, including how conducive the setting is to co-worker communication about pay (length of time together, physical proximity, privacy). We embed the full job description from TaskRabbit in a series of vignettes on Amazon's Mechanical Turk, another online labor platform. Each vignette describes two people who have placed private bids for the job, and meet each other for the first time upon commencing the work. The story is followed with five questions; How much time will these two workers be working together? How likely is it that they will discover what the other one is being paid for the same work? How likely is it that one could be a leader or role model and influence the work of the other? How likely is it that the employer sees the effort that each of them individually contributed? and How likely is it that the same employer would hire them in the future for a similar purpose? The average inter-rater correlation coefficient is 0.37 in the reported likelihood of learning the bid of their co-worker.[21] In Table 1 we report characteristics of workers, employers and task by whether the job has a publicly posted price.

4.4 Verifying Bargaining Assumptions

We based our theoretical model on the (re)bargaining process on TaskRabbit, namely a worker can propose a take-it-or-leave it offer to her employer on the job. The premise of our model has two clear empirical implications for the outcome of a re-bargaining process in multi-worker, co-located tasks:

SF1: Workers are no more or less likely to receive a higher wage based on how far their bid is from the highest accepted bid.

SF2: Workers who receive different wages than their initial bids receive a wage equal to the highest accepted bid.

We use TaskRabbit data to verify these two stylized facts, lending credence to our modeling decisions, which we show in Table 3. Conditional on renegotiating a higher wage than the initial contract, similar workers negotiate pay nearly equal to that of the highest accepted bidder (Col. 6, SF1). At the same time, the distance between bidders adds no additional predictive power as to whether or not a worker receives any raise at all (Col. 1-3, SF2).

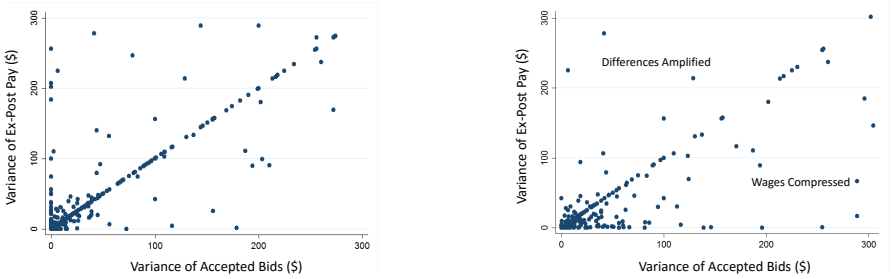
In the following section, we show a relationship between the communication channel between co-workers and the likelihood that wages are equalized. We present empirical tests that renegotiation is the cause of this wage compression, as opposed to alternative mechanisms such an productivity spillovers or employer preferences for equity. In Appendix Section A we go into more detail about ruling out these alternatives.

[21] We find statistically significant differences in the reported likelihood of communicating pay when we vary the gender of the names of the workers in the vignette. We use gender-specific measures in an alternative specification.

4.5 Quantifying Wage Compression

Theorem 1 states that increased transparency leads to compression in wages of employed workers. We first present a visual depiction of wage compression across two types of tasks, higher transparency jobs that require multiple workers to be situated together and lower transparency jobs that can be carried out separately. For example, packaging items at the office under a tight deadline as opposed to distributing marketing material in different districts.

Figure 1: Variance in final pay vs. bids



Notes: Each observation summarizes the variance in pay among the workers that have been selected for multi-worker separated jobs (Panel a) and multi-worker co-located jobs (Panel b). The x-axis is the variance in the bids accepted for a job, in dollars. The y-axis is the variance in the final payout. An observation below the 45 degree line indicates that wages are compressed on the job, while and observation above the 45 degree line indicates that wages become dispersed on the job.

Figure 1 offers a visual depiction of the variance in wages for workers assigned to the same job on TaskRabbit. For each multi-worker task, we plot the variance in the bids that have been accepted (x-axis) against the variance of the final payouts (y-axis). For those observations along the 45 degree line, variation in bids and payout remain constant, as would be the case if the employer makes no adjustment to the initial contract. Observations that fall along the x-axis represent tasks which received unequal bids, but for which final pay was equal for all workers. The concentration of observations that fall beneath the 45 degree line illustrate the tendency for employers to reduce the variance in ex-post payment relative to the variance of ex-ante bids.

Among co-located workers, 27% receive pay that is higher than their bids, as opposed to 10% on average when workers are separated. The Gini coeffient of final pay is, on average, two-thirds the Gini coefficient of selected offers when workers are co-located, and cannot be statistically differentiated from 0 when separated (Table 4).

In a regression at the level of an individual worker, we demonstrate a co-workers' bid impacts own final pay with individual productivity measures included as additional co-variates. To interpret co-workers' bid as having a causal effect on own final pay when workers are co-located, bidding cannot be strategic nor can employer's selection of workers as a function of

co-location. In Appendix Section A.2 we provide empirical support for these assumptions, using the visibility of multiple workers in the job posting and variation in the supply of bids received by the employer for identification, along with a discussion of why we do not observe more strategic bidding.

Our measure of individual productivity uses the lifetime review of a worker on the platform.[22]

We run the specification in Equation 9 below. Each accepted bid placed by worker i is one observation. The subscript s refers to the job and j to the employer.[23] The dependent variable is the difference between ex-post payment and ex-ante bid, Δy_{ijs}, expressed as a percentage raise above a worker's initial contract bid. Distance between a worker's initial bid and that of the highest selected bidder is also expressed as the percentage above initial bid contract, T_{ijs}.[24] C denotes the set of job settings. We differentiate between three types of multi-worker settings, we define a subset $\{\theta_1, \theta_2, \theta_3\} \subset C$, corresponding to co-located jobs, segregated local jobs and virtual jobs respectively. For our main specification we examine local jobs that require workers to be together or separate.

$$\begin{aligned} \Delta y_{ijs} = \alpha_0 + \alpha_j + \beta \bar{X}_i + \phi \bar{X}_s + \epsilon_{ijs} \\ + \gamma_{\theta_1} T_{ijs} + \gamma_{\theta_2} T_{ijt} \mathbb{1}_{Separate} \end{aligned} \tag{9}$$

These results can be seen in Table 9. In Col. 1 through 3, the coefficient γ_{θ_1}, indicates that when workers are co-located, an additional 10% gap between initial bids and the high bidder will result in a 4% increase in ex-post pay. The effect of the distance between co-worker bids on the final pay when workers are separated physically, $\gamma_{\theta_2} + \gamma_{\theta_1}$, cannot be statistically distinguished from 0. In each column moving right to left with include more controls. Col. 3 demonstrates this finding is robust within employer.

[22]The literature on user generated content has identified a number biases and manipulation techniques that we can address using data about performance that the platform collects but is not visible to the users. Nosko and Tadelis (2015) show the "sound of silence," or missing reviews, on Ebay is skewed toward negative feedback. We show the share of missing reviews on TaskRabbit predicts whether an employer returns to the platform, TaskRabbit's central measure of employer satisfaction, and the star-rating conditional on receiving a rating (Table 12 Col. 1 and Col. 2 respectively). Another important feature of this unobserved measure is that it is correlated with ex-post pay, but not the ex-ante bid accepted (Table 13), suggesting that we are really detecting the performance that the employer observes on-the-job.

[23]When we include employer fixed effects, we also include many characteristics of the task itself in our analysis.

[24]We are confident that there is no measurement error in the bids, and are therefore comfortable normalizing both the dependent and independent variable by the initial bid.

4.6 Employment

TaskRabbit administrative data includes those job posts that expired before they matched, offering us a measure of unmet labor demand.[25] We refer to the employment rate, in the context of TaskRabbit, as the proportion of positions filled.

Theorem 2 finds that transparency raises employment by more when the value of labor to the employer is low. To test this finding, we need to measure both employment and employers' value of labor.

We do not directly observe the maximum an employer is willing to pay to fill the job vacancy. We do however observe self-reported annual earnings from household employers.[26] We use annual earnings as a proxy for willingness to pay. One reason to favor this measure is that, from a survey of employers conducted by TaskRabbit, the most common alternative to using TaskRabbit to complete a task is to do it oneself. Using money as measure of the opportunity cost of time, higher income employers are more likely to have higher time costs (i.e. leisure is a normal good).[27]

We find that below-median earners (under $150,000 annual income) choose to use a transparent, posted price scheme more often than high earners (5% more often, Table 5). The gap in the employment rate between transparent posted price jobs and private auctions is also slightly larger for low earners even before accounting for the selection which dampens this effect (3% larger, see Table 6). Overall employment is increased by 12% under transparent posted prices. We reproduce this analysis restricting the sample to the first three jobs posted by each employer to minimize the effects of learning and strategic selection and find similar results.

4.7 Profit Share

Theorem 3 predicts that higher levels of transparency are associated with higher employer profits on average. We note in Table 10 that both bid and final pay are approximately 10% lower when a price range is mentioned in a job post compared to those with no mention, a result that is robust across specifications with worker fixed effects and job characteristic controls. Combined with insignificant differences in close rates, lower wages likely translate directly into higher profits.[28] We find similar effects of transparency on profits if we compare

[25] Cullen and Farronato (2016) estimate a slack labor market and highly elastic labor supply, supporting this notion that unfilled tasks reduce total work completed by TaskRabbit workers and wages paid to TaskRabbit workers.

[26] TaskRabbit also hires third party companies to report socioeconomic characteristics of employers using a combination of address, job title, and other public records.

[27] We directly manipulate the value of labor to the employer in our field experiment which is described in Section 5.

[28] We directly observe employer profits in our field experiment which is described in Section 5.

posted price jobs (higher transparency) to private auctions (lower transparency), and if we compare co-located, multi-worker jobs (higher transparency) to separated, multi-worker jobs (lower transparency).

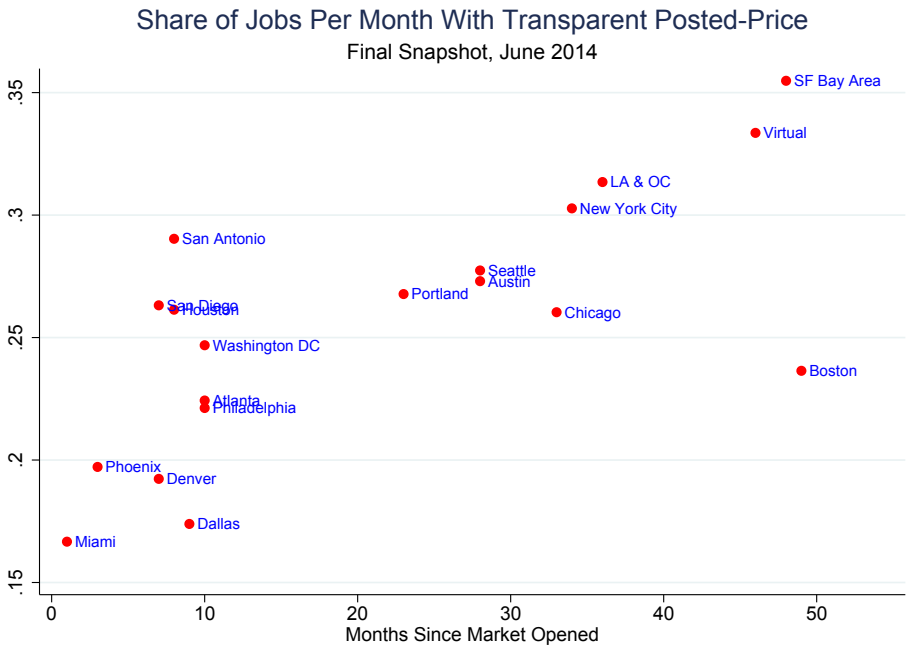
4.8 Endogenous Choice of Transparency

In Corollary 2 we find that there is not an equilibrium in which the firm selects either full transparency or full secrecy depending on its value of labor, as such a scheme unravels toward all firm types pooling on full transparency. Employers make exactly this choice for single worker, and multi-worker, separated jobs: if the employer elects to select from worker bids, there will be no transparency, and if the employer elects to post a price, there will be full transparency. We observe the predicted unraveling in TaskRabbit. All else equal, the proportion of posted price jobs rises by 1% per month, which we show in Table 7. Observed unraveling in labor markets typically takes time.[29] This is consistent with a bounded rationality explanation a la Kandori et al. (1993) in which employers select to post a price or accept bids based on which scheme would have maximized their expected profit in a previous period. [30]

[29] Roth and Xing (1994) discuss this timeline in a number of labor markets.

[30] Employers need not even rely on their own experiences as TaskRabbit pricing discussion threads exist on websites including Glassdoor, Quora, and Reddit, in addition to word-of-mouth information acquisition. Furthermore, a website including empirical analyses of "optimal" TaskRabbit pricing exists, meaning that employers can potentially use the data of previous job posters to optimally select their pricing strategies. For example, Kerzner (2013) contains publicly available empirical analyses of pricing strategies in TaskRabbit.

Figure 2: Posted price and market age



Notes: Figure 2 plots the age of each TaskRabbit market (horizontal axis) and the proportion of posted price jobs in each market (vertical axis) at the end of our data sample in June, 2014. Older markets appear to be associated with a higher proportion of posted price jobs. "Virtual" refers to tasks that are completed by workers online. As location is not relevant for these types of markets, TaskRabbit rolled out virtual tasks country-wide at the same time. TaskRabbit entered Boston in 2008 nearly two years before the start of our data sample. In our analysis we treat the Boston market as if it started at the same date as our data sample, but in reality, there are many observations that we do not observe.

Figure 2 shows the share of posted price jobs in each TaskRabbit market in June of 2014, in which older markets are generally associated with a higher proportion of posted price jobs. In June, 2014 TaskRabbit removed the bid acceptance procedure from all markets. We discuss alternative hypotheses for this market trend, and why we do not believe these are plausible, in Appendix A.3.

5 Study II: Evidence from Field Experiment

We conduct a field experiment to further test our findings.[31] We hire 347 "managers" and 1047 "workers" from an online labor market who are tasked with negotiating wages for a real-effort task in which we exogenously vary transparency. The experiment relies on free-form bargaining between workers and managers. We view this as a strength of our experiment; by allowing for more natural bargaining protocols than the TaskRabbit environment, we take

[31] Additional details, including the code used to run the experiment, can be found on our academic websites.

another step toward applying our findings to traditional labor markets.

We directly measure worker cost of effort, productivity, outside option and employer profits. This allows us to explore additional notions of fairness, such as compression in worker surplus in addition to compression in earnings. In Section 6 we use this additional data to test the relative importance social concerns surrounding pay transparency, and run a horse race between such a model and our re-bargaining model.

5.1 Procedure

Participants are recruited from Amazon's Mechanical Turk between October, 2016 and May, 2017. Participants are assigned to either the role of "worker" or "manager" and informed that their participation is voluntary and part of an academic experiment. All participants are given the following instructions: managers and workers are tasked with negotiating a per-page rate for completing text-to-text transcription of US Census tables from the 1940s. If a worker and manager agree on a wage and the worker completes a page above a stated accuracy threshold, the worker receives the agreed upon wage. Workers are each able to complete up to 5 pages of transcription. Each manager is assigned to an average of 4 workers and privately given a budget amount. Managers are paid this budget, minus payments made to the worker, for every completed page above the accuracy threshold. Therefore, there is no direct impact of worker turnout on marginal manager earnings.

Before interacting, workers are shown a sample transcription page. They are asked to place a per-page bid for completing a similar transcription, and this bid is shared with the manager. We also collect data about the minimum price a worker would accept to transcribe a page, and make it clear to workers that this information will not be shared with managers. Similarly to the method Becker et al. (1964), we make truth-telling a dominant strategy for workers; they are asked to make several selections between receiving \$X for completing one transcription page or \$2 for doing nothing. We vary X and randomly select one choice for 1 in 10 workers and give the worker their reward (either \$2 or the opportunity to complete an additional page at \$X) after their initial assignment has been completed. We also survey other characteristics including the expected time it takes to complete each page, daily household income, management/transcription experience, gender, location, and age.

Upon answering intake questions, managers are shown worker bids. Workers and managers meet and bargain for wages in anonymous online chatrooms. No other modes of communication occur between participants. We place no restrictions on the way in which participants bargain, only that they indicate in the chatroom a final agreed upon wage. Participants are told that we can monitor chatrooms and that agreement is required to dispense payment. After agreeing to a wage, each worker has 48 hours to complete up to 5 pages of transcription. All pages submitted within the time frame meeting our accuracy threshold

result in payment to worker and manager.

5.2 Treatments

The experiment follows a 2×2×2 design. One experimental condition is the public visibility of wage negotiations. Managers either negotiate wages with each worker in a separate, private chat room ("secrecy" treatment), or the manager negotiates over a common chat room ("transparency" treatment) with all workers. A second experimental condition varies the budget assigned to managers, either \$5 per page or \$9 per page. A third treatment either requires managers to accept all bids less than or equal to their budget, or allows them to actively bargain with workers.

5.3 Administrative details

We present results from two procedures that differ only in the degree of automation. One version relies on us, the experimenters, to invite workers to chat rooms and collect transcriptions via email following the intake survey, and another is completely automated in this dimension, as all interaction occurs through a single web interface programmed in oTree (Chen et al., 2016). oTree became available after the initial rounds of our experiment. More than 3/4 of our participants (800 of 1047 workers and 253 of 347 managers) interacted through the automated system, and results are comparable across the two interfaces.

The mean bid placed by workers is \$9.50 per page, and roughly one-third transact at \$5 or less per page. Table 2 shows that, along ex-ante observable characteristics, the workers and managers randomly assigned to different treatments are comparable.

Transcription accuracy is calculated using the Levenshtein distance measure (Levenshtein, 1966), defined as the minimum number of single-character edits (substitutions, deletions, or insertions) necessary to change one string into another. Each submitted page with a Levenshtein distance to the original document of fewer than 5% of the total number of characters on the page meets our accuracy threshold.

5.4 Analysis and main findings

Workers often used the wages of others to bargain for wages in the transparency treatment. Below, we provide a portion of a wage negotiation.

Example dialogue:

Manager: You agreed to \$1 per page?

Worker: I really don't remember, it sounds good but I suppose you would give the same to everyone? I see you gave 5$ to [OTHER WORKER]

Manager: You said you have experience, so...I have a question related to changing your fee :). Can you do at least 95% accurate work?

Worker: off course I can! I will triple check to be sure too!

Manager: Okay, $5 per page!

Wage and Surplus Compression:

Pay is significantly compressed in the transparency treatment compared to the secrecy treatment, even when we include workers who do not complete the task in our analysis. In every instance where the worker reached an agreement with the employer in the transparency treatment, the worker ultimately received pay equal to the highest paid worker (See Table 4). This provides additional evidence in support of theoretical predictions SF1 and SF2. The similarity in findings with TaskRabbit data is powerful. In the experimental setting we can be certain that there are no productivity spillovers between workers.

We present relative worker surplus as an additional measure of fairness, which should be interpreted in light of what outside option value represents. If the outside option reflects cost of effort, then compression in the piece rate shifts worker surplus toward the person most capable of completing the task quickly. If the outside option is simply a function of previous wages, then compressing wages may lessen disparities in outside opportunities.

We measure worker surplus as the final payout received less the reservation value that we elicited on a worker intake survey using the BDM procedure introduced in Becker et al. (1964). On average, the Gini measure of dispersion for worker surplus is 0.115in the secrecy treatment, and more than *doubles* in the transparency treatment (Table 4). Is this evidence of greater inequality in welfare?

We find evidence that outside options reflect cost of effort, in part, by eliciting the time it takes to complete a page of transcription. Indeed, rising dispersion in worker surplus reflects a shift of surplus towards those with low cost of effort and away from those with high cost of effort. When we convert the piece-rate contracts into an hourly wage, we no longer find strong evidence of wage compression (Table 4).

Employment:

Our definition of employment in this context is the match rate, the proportion of workers who agree on a wage with their manager.

Recall Theorem 2 states that the benefit of transparency to employment is higher when the value of labor to the employer is low. This experiment provides a direct test of the theorem.[32] We show in Table 6 that when the value of labor is \$5 per unit of output, employment rises by 20% when negotiations move from the secret to transparent forum. On the hand, employment falls by 10% when the manager's budget is \$9 as we move from secrecy to transparency. This is evidence that higher transparency is more effective at improving employment when the employer value (manager budget) is low.

Profits:

Our theory predicts that, irrespective of the value of labor to the employer, employer profits will be higher in expectation under full transparency relative to full secrecy. In Table 11 we display empirical evidence consistent with this prediction. Manager profits are over 50% in the transparency group than the secret group. In private forums, workers bargain more aggressively and negotiations more often result in disagreement and no transaction.

6 Alternative model

We investigate an alternative channel, fairness concerns, that has been suggested in the literature as a possible explanation for compressed wages in pay-transparent environments. Several papers (Card et al. (2012), Mas (2016), Perez-Truglia (2016), and Breza et al. (2016)) document a morale effect when workers learn they are underpaid relative to their peers; upon learning that she is underpaid, a worker is less likely to put forth high effort at work. This notion is based on the fair wage-effort hypothesis (Akerlof and Yellen, 1990) which posits that workers face an additional cost to effort upon learning they are underpaid.[33] A proactive employer may augment the wages of workers who learn they are underpaid in order to avoid low effort provision.[34] Indeed, this is consistent with a world in which the firm has all the bargaining power in wage re-negotiations.[35]

But is this alternative model consistent with our empirical results?

[32] In TaskRabbit we relied on characteristics of the employer to proxy for the value of labor.

[33] Alternatively, the modeling in this section can be interpreted as a social, "fairness" cost that an employer pays for having known, unequal wages.

[34] This, of course, hinges on the ability of the firm to monitor a worker's output or knowledge of the wages of others within the firm.

[35] Suppose, for example, that firms could make TIOLI offers to workers after observing their outside options. In equilibrium, the firm would select a subset of workers to employ and pay each worker her outside option. Once a worker learns the wages of her peers, the firm would (possibly) increase her wage to offset morale costs.

We demonstrate that the only way that morale can account for the stylized facts we have presented is if a worker puts in 0 effort when she learns she is paid any ϵ less than another worker. We see no evidence of such an extreme morale response in our experiment. When the manager does not equalize wages in the transparency treatment, completion rate falls by (only) 20 percent. Breza et al. (2016) report a reduction in output of about one-third when workers learn that a co-worker earns 5% more.

To derive our test of competing mechanisms, we build a model of endogenous worker effort. Suppose that workers make initial offers w_i^* and the wage profile of the firm arrives at rate λ independently to each worker. Additionally, let each worker i select $e_{i,t} \in [0,1]$ which is the probability of successfully completing her time t duties and getting paid her time t wage $w_{i,t}$. We re-imagine θ_i as a cost of effort. All workers have the same outside option, normalized to 0 and have to pay a linear flow cost of $\theta_i \cdot e_{i,t}$ to put in effort. We focus our attention exclusively on "myopic" equilibria in which $e_{i,t}$ is chosen to maximize time t profit for each worker.[36]

First, we show that when workers hold the bargaining power and can make TIOLI offers to the firm at any time, this alternative model leads to similar equilibria as the base model in the paper; any equilibrium of the original model has a counterpart in the endogenous worker effort model in which all employed workers select $e_{i,t} = 1$ for all periods during which they are employed.

PROPOSITION **4.** *For every equilibrium of the original game satisfying A1-5, there exists an equilibrium of the endogenous worker effort game in which all employed workers set $e_{i,t} = 1$ for all periods of employment and all other actions are the same on equilibrium path by all agents.*

We have seen that the bargaining mechanism is unchanged in the presence of endogenous worker effort of this form, so all of the other results of the paper carry through to this setting. We now remove bargaining power from workers, that is, workers can make a wage offer w_i^* when they are initially hired, and thereafter can only make effort choices. Instead, the firm can observe whether each worker has received wage information at each time t before deciding whether or not to unilaterally increase the wage of each worker. We refer to this as the proactive employer model. In order to give the firm a reason to increase the wages of workers, we include a morale cost $m(e,d)$ where $d = \bar{w} - w_{i,t}$. We assume $m(\cdot,\cdot)$ is non-decreasing in both arguments and that $m(0,\cdot) = m(\cdot,0) = 0$. Suppressing time and worker identity notation, before learning about the wages of her peers a worker's payoff is $w \cdot e - \theta \cdot e$. As before,$w \geq \theta$ and due to the linearity of costs, a worker will put in full effort in equilibrium. Upon seeing the wages of her co-workers and learning $\bar{w}$, the worker's flow

[36]This is to rule out equilibria that stem from non-credible threats, such as workers electing to exert low effort if the firm does not raise their wage at an arbitrary time.

payoff becomes $w \cdot e - \theta \cdot e - m(e, d)$. Depending on $m(e, d)$, the worker may optimally shirk if wages are augmented. We now formally state conditions on the morale function for the proactive employer model to fit the two stylized facts.

PROPOSITION **5.**

1. *The firm will increase the wage of a worker i at time t only if i learns the wages of her co-workers at time t.*

2. *An equilibrium of the original model has a counterpart in the proactive employer model in which the firm sets $w_{i,t} = \bar{w}$ for the duration of any worker's tenure at the firm upon her learning $\bar{w}$ for every $\Lambda, v, r,$ and s if and only if $w \cdot e - \theta \cdot e - m(e, d) \leq 0$ for any $e \in [0, 1]$ and any $d \in (0, 1]$.*

This result states that only very extreme morale cost functions give equivalent predictions as the bargaining model; unless every worker would optimally choose to quit her job (put in 0 effort) upon receiving any wage less than $\bar{w}$, the firm will not always equalize the wages of all workers. The intuition is that when transparency is low, firms make close to zero profit from their highest paid worker ($v - \bar{w} \approx 0$) so even if a worker drastically reduces her effort, but does not quit upon learning she is underpaid, a proactive firm would still prefer not to increase her wage all the way to $\bar{w}$. Practically, the severity of this result implies that full wage equalization should not be expected for "smooth" morale cost functions. If transparency is unanticipated, as we discuss in Appendix A.2, then it becomes even more implausible that morale is the cause for wage equalization; if workers and the firm place and accept initial offers thinking there is full secrecy, then $\bar{w} = v$. But then for any actual level of transparency $\Lambda \in (0, 1)$ full wage equalization in line with SF1 and SF2 would lead the firm to derive zero profits from any worker.

There is evidence to suggest that there is a morale cost to effort when a worker learns she is underpaid however, most workers complete their task.[37] The transcription rate in our experiment drops by 18% in then treatment in which managers are unable to negotiate away pay differences and negotiations take place in a transparent, common chat room compared to the case in which managers are allowed to negotiate in a transparent, common chat room..[38] However, the task completion rate when managers where unable to negotiate with workers workers was still statistically different from 0 at the .01 level. Furthermore, the transcription rate falls the farther a worker's bid is from the highest accepted bidder in this treatment, as

[37]Robert Duvall famously refused to take part in The Godfather Part III stating, "if they paid Pacino twice what they paid me, that's fine, but not three or four times, which is what they did." http://www.imdb.com/name/nm0000380/bio.

[38]As managers accepted all bids below the budget in this treatment, whereas managers often (optimally) used a common wage strictly less than their budget when they were allowed to negotiate, we believe this to be a lower bound on the effect of morale.

we show this in Figure 6. There are similar findings in the literature (Card et al. (2012), Mas (2016), and Breza et al. (2016)). Therefore, in light of Proposition 5, and SF1 and SF2, we find that the mechanism equalizing wages in co-located, multi-worker jobs on TaskRabbit is more in line with re-bargaining than a proactive employer optimally responding to morale costs of workers. We reiterate that we find evidence supporting the existence of a morale cost to learning that one is underpaid, but in the presence of renegotiation, this does not change our findings, as workers will bargain away wage differences and therefore face no lingering morale costs.[39]

7 Gender differences and the gender pay gap

We allow for differences between workers along two dimensions, outside options and heterogeneities within the communication network. We extend our empirical and theoretical analysis to shed light on how transparency affects wage inequality when genders differ along these dimensions. We believe these to be important avenues of inquiry, as pay transparency is commonly cited as a way to close the gender wage gap, which stands at roughly 8% after controlling for observables (Blau and Kahn, 2016; Goldin, 2014).

We elicit the outside option of workers in our experimental setting as discussed in Section 5. We find that women have on average 9.7% lower outside options than men, and we show that increasing pay transparency can close a pay gap caused by differences in outside options.

To see this theoretically, let there be two types of workers m (male) and f (female) such that $qG_m(x) + (1-q)G_f(x) = G(x)$ for all $x \in [0,1]$, where $G_m(x)$ and $G_f(x)$ are both atomless distributions and $q \in [0,1]$ is the proportion of men in the market. The first, and simplest result, is an application of Theorem 1. Similarly to above, we denote the average equilibrium lifetime earnings of an employed worker of type $\ell \in \{m, f\}$ as $T(\Lambda, v, \theta_i, \ell)$.

COROLLARY **3.** *If $G_m(\cdot)$ first-order stochastically dominates $G_f(\cdot)$ then $\frac{\mathbb{E}_{G_f}[T(\Lambda,v,\theta_i,f)]}{\mathbb{E}_{G_m}[T(\Lambda,v,\theta_i,m)]}$ converges monotonically to 1 as Λ converges to 1 for all v.*

In words, this result says that the average earnings of employed women is rising relative to the average earnings of employed men as transparency increases, and reaching full transparency completely equalizes earnings. The proof of this result follows from Theorem 1. When $G_m(\cdot)$ first-order stochastically dominates $G_f(\cdot)$, it is possible to pair up every f type worker with an m type worker with higher outside option. Formally, let $\mu : [0,1] \to [0,1]$ define for each f type worker i an m type worker j such that $\theta_j \geq \theta_i$ and $\mu(\theta_i) \neq \mu(\theta_{i'})$ for

[39] If we use the morale specification in this section and give the workers the ability to make TIOLI offers to the firm, then workers optimally request $\bar{w}$ upon seeing this value as it improves their objective function without affecting their constraints regarding chosen effort. This means that on path, workers will renegotiate successfully to $\bar{w}$ and never pay the morale cost or put in low effort.

any $i \neq i'$. We know by Theorem 1 that $\frac{T(\Lambda,v,\theta_i)}{T(\Lambda,v,\theta_j)}$ converges monotonically to 1 in Λ, which implies that the average income of each worker type also converges monotonically to 1 in Λ.

However, if transparency relies on word-of-mouth, and co-worker networks are dominated by one gender or the other, permitting communication exacerbates the gender pay gap. Empirically, we find evidence that network effects heavily mediate the impact of partial transparency, or open communication, on the gender pay gap. First, we find that men receive bonuses more often than women on average when workers are co-located, and not when they are separated (Table 8). Second, when the majority of workers at a job are female, all workers, male and female, receive bonuses less often doing co-located work, while these composition differences do not affect the outcome when workers are separated. This is in line with findings of our worker surveys, where men report a higher likelihood of learning about co-worker pay on the job (Figure 7).

We make simple adjustments to the model to capture these network effects. Let $\alpha_m > \alpha_f > 0$ be the rates at which men and women receive wage information, respectively, and let q be the proportion of men in the industry. Let the arrival rate of information for a worker of gender $\ell \in \{m, f\}$ be $\alpha_\ell q \lambda$. Then

$$\Lambda_\ell = \frac{\alpha_\ell q \lambda}{\rho+\delta+\alpha_\ell q \lambda} \quad \text{for } \ell \in \{m, f\} \quad \text{and } \lambda \in [0, \infty) \tag{10}$$

We plot the difference in arrival rates between men and women, $\Lambda_m - \Lambda_f$, as a function of λ for arbitrary parameters in Figure 5. $\Lambda_m - \Lambda_f$ is initially decreasing in λ but over time decreases toward 0.

We find that this shape holds for all parameters. Network effects cause the de facto arrival rate of information of men to be greater than that of women. It is easy to see that for all $q > 0$ and $\lambda \in (0, \infty)$, $\Lambda_m - \Lambda_f > 0$. If $q = 0$ or $\lambda \in \{0, \infty\}$ then $\Lambda_m - \Lambda_f = 0$. The following result shows that increasing transparency below a certain threshold increases the difference in the arrival rate of information for men and women.

PROPOSITION **6.** *Let λ_c solve $\frac{\alpha_m}{\alpha_f} = \frac{(\rho+\delta+\alpha_m q\lambda)^2}{\left(\rho+\delta+\alpha_f q\lambda\right)^2}$. $\Lambda_m - \Lambda_f$ is strictly increasing in λ for all $\lambda < \lambda_c$ and strictly decreasing for all $\lambda > \lambda_c$.*

Compare the effects of moving from zero transparency to some $\lambda > 0$. When λ is relatively low, information transmission between workers happens rarely through word of mouth. Because men are more likely to speak about wages, they disproportionately benefit from low levels of transparency compared to women. However, when λ is relatively high, men gain relatively less compared to women because all workers learn information quickly regardless of their communication propensity. In extreme cases of transparency in which the firm posts a price ($\lambda = \infty$) any communication advantage men have completely disappears as information arrival is not based on word of mouth transmission.

An important implication of these findings is that intermediate levels of transparency, or transparency that relies on communication channels between co-workers, may exacerbate pay discrepancies while full transparency reduces them. Advocates of gender equality may reconsider whether pursuing and penalizing employers that prohibit pay discussions will further their cause.

8 Conclusion

Although pay transparency has been in the political and popular spotlights, the equilibrium effects of making pay more transparent have not received adequate attention. First, does it result in more equal wages for similarly productive workers? Second, does it lead to higher employment rates? Third, how does it affect the division of surplus between firms and workers?

We investigate these questions theoretically and empirically by introducing a model in which workers optimally negotiate salaries in the presence of dynamic arrival of wage information, and testing the model's predictions using big data from TaskRabbit, and a field experiment.

We find, theoretically and empirically, that increasing pay transparency can increase employment and decrease inequality. We also find that increasing pay transparency may shift surplus away from workers and toward firms. Moreover, intermediate levels of pay transparency, achieved through a permissive environment to discuss relative pay, can exacerbate the gender pay gap by virtue of how information spreads.

There may be a direct need for government intervention in order to maintain a desirable level of transparency. We have shown through our model that any scheme in which firms choose transparency based on private characteristics is unsustainable, as the signal sent to prospective workers is sufficiently strong to cause unraveling toward full transparency. We observe this unraveling among employers on the TaskRabbit platform, and more generally among startups in Silicon Valley that are increasingly opting to be fully transparency about pay.

Our analysis rests on a methodologically diverse approach, which we view as a strength of this paper. Our simple, equilibrium model of pay transparency reveals consequences of pay transparency that were not initially apparent to us, and allow us to further the literature on this policy. With TaskRabbit administrative data, we have the rare opportunity to observe an employer's active decision to adjust pay across different settings where large pay disparities arise through the bidding process for work. In many ways, this data is truly ideal: we observe initial and renegotiated wages, there is rich demographic information on both workers and employers, and we have the ability to view jobs that are not completed.

The experiment we run on Amazon's Mechanical Turk allows for estimation of unobservables in TaskRabbit, and verification of our findings in a controlled data setting. Most importantly, the simplified bargaining protocol used in TaskRabbit and built into our model is replaced by free-form bargaining between participants. Our experimental findings match the findings in TaskRabbit and our model, providing a basis to believe our conclusions regarding the effects of pay transparency policies may extend to other labor markets.

References

John M. Abowd, Francis Kramarz, and David N. Margolis. High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333, 1999. 1.1

George A. Akerlof and Janet L. Yellen. The Fair Wage-Effort Hypothesis and Unemployment. *The Quarterly Journal of Economics*, 105(2):255–283, 1990. 1.1, 6

Linda Babcock and Sara Laschever. *Women Don't Ask: Negotiation and the Gender Divide*. Princeton University Press, 2003. 1.1

Jesper Bagger, François Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics. *American Economic Review*, 104(6):1551–1596, 2014. 1.1

Gordon M. Becker, Morris H. DeGroot, and Jacob Marschak. Measuring Utility by a Single-Response Sequential Method. *Behavioral Sciences*, 9(3):226–232, 1964. 1, 5.1, 5.4, 4

Truman Bewley. *Why Wages Don't Fall During a Recession*. Harvard University Press, 1999. 1.1

Francine D. Blau and Lawrence M. Kahn. The Gender Wage Gap: Extent, Trends, and Explanations. *NBER Working Paper 21913*, January 2016. 1, 1.1, 7

Francis Bloch and Harl Ryder. Two-Sided Search, Marriages, and Matchmakers. *International Economic Review*, 41(1):93–115, 2000. 47

Emily Breza, Supreet Kaur, and Yogita Shamdasani. The Morale Effects of Pay Inequality. *mimeo*, 2016. 1, 1.1, 6, 6

Kenneth Burdett and Melvyn G. Coles. Long-Term Partnership Formation: Marriage and Employment. *The Economic Journal*, 109(456):F307–F334, 1999. 47

Kenneth Burdett and Dale T. Mortensen. Wage Differentials, Employer Size, and Unemployment. *International Economic Review*, 39(2):257–273, 1998. 1.1

Pierre Cahuc, Fabien Postel-Vinay, and Jean-Marc Robin. Wage Bargaining with On-the-Job Search: Theory and Evidence. *Econometrica*, 74(2):323–364, 2006. 1.1

David Card, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. Inequality at Work: The Effect of Peer Salaries on Job Satisfaction. *American Economic Review*, 102(6):2981–3003, 2012. 1.1, 6, 6

Hector Chade. Matching with Noise and the Acceptance Curse. *Journal of Economic Theory*, 129 (1):81–113, 2006. 47

Gary Charness and Peter Kuhn. Do Co-Workers' Wages Matter? Theory and Evidence on Wage Secrecy, Wage Compression and Effort. *Discussion Paper Series, IZA DP No. 1417*, 2004. 1.1

Kalyan Chatterjee and William Samuelson. Bargaining under Incomplete Information. *Operations Research*, 31(5):835–851, 1983. 1.1, 2.4

Alice Chen, Emily Oster, and Heidi Williams. Why is Infant Mortality Higher in the US than in Europe? *Mimeo*, 2014. 1

Daniel L. Chen, Martin Schonger, and Chris Wickens. oTree-An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97, 2016. 5.3

Zoe Cullen and Chiara Farronato. Outsourcing Tasks Online: Matching Supply and Demand on Peer-to-Peer Internet Platforms. *mimeo*, August 2016. 1.2, 25, 8

Alex Edmans, Xavier Gabaix, Tomasz Sadzik, and Yuliy Sannikov. Dynamic CEO Compensation. *Journal of Finance, Forthcoming*, 2016. 12

Liran Einav, Chiara Farronato, Jonathan D. Levin, and Neel Sundaresan. Sales Mechanisms in Online Markets: What Happened to Internet Auctions? *mimeo*, 2013. A.3

Diana Farrell and Fiona Greig. Paychecks, Paydays and the Online Platform Economy. *JPMorgan Chase & Co. Institute*, 2016. 1.2

Robert Frank. Are Workers Paid Their Marginal Products? *American Economic Review*, 74(4): 549–571, 1984. 1.1

Xavier Gabaix. Power Laws in Economics and Finance. *Annual Review of Economics*, 1(1):255–294, 2009. 12

Xavier Gabaix. Power Laws in Economics: An Introduction. *Journal of Economic Perspectives*, 30 (1):185–206, 2016.

Claudia Goldin. A Grand Gender Convergence: Its Last Chapter. *American Economic Review*, 104(4):1091–1119, 2014. 1, 1.1, 7

Marina Halac. Relational Contracts and the Value of Relationships. *American Economic Review*, 102(2):750–779, 2012. 48

Robert Hall and Alan Krueger. Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-Job Search. *American Economic Journal: Macroeconomics*, 4(4):56–67, October 2012. 3, 17

Michihiro Kandori, George J. Mailath, and Rafael Robb. Learning, Mutation, and Long Run Equilibria in Games. *Econometrica*, 61(1):29–56, January 1993. 4.8

Lawrence D. Katz and David H. Autor. Changes in the Wage Structure and Earnings Inequality. In Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*. 1999. 1

Lisa Keister. The One Percent. *Annual Review of Sociology*, 40:347–367, July 2014. 1.1

Elena Kerzner. Lessons Learned from Creating my TaskRabbit Pricing Strategy. *http://www.priceintelligently.com/blog/bid/192053/Lessons-Learned-from-Creating-My-TaskRabbit-Pricing-Strategy*, 2013. 30

Wolfgang Leininger, Peter B. Linhart, and Roy Radner. Equilibria of the Sealed-Bid Mechanism for Bargaining with Incomplete Information. *Journal of Economic Theory*, 48:63–106, 1989. 1.1, 2.2

Vladimir I. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10(8):707–710, 1966. 5.3

Matt Marx. The Firm Strikes Back Non-Compete Agreements and the Mobility of Technical Professionals. *American Sociological Review*, 76(5):695–712, 2011. 16

Alexandre Mas. Does Transparency Lead to Pay Compression? *Journal of Political Economy, Forthcoming*, 2016. 1.1, 6, 6

Roger Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981. 3.3, D.1, D.1

Muriel Niederle and Lise Vesterlund. Do Women Shy Away From Competition? Do Men Compete Too Much? *Quarterly Jounal of Economics*, 122(3):1067–1101, 2007. 1.1

Chris Nosko and Steven Tadelis. The limits of reputation in platform markets: An empirical analysis and field experiment. *Working paper*, 2015. 22, A

Ricardo Perez-Truglia. The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment. *mimeo*, 2016. 1.1, 6

Thomas Piketty. *Capital in the Twenty-First Century*. Harvard University Press, 2014. 1

Fabien Postel-Vinay and Jean-Marc Robin. The Distribution of Earnings in an Equilibrium Search Model with State-Dependent Offers and Counteroffers. *International Economic Review*, 43(4): 989–1016, 2002. 1.1

Fabien Postel-Vinay and Jean-Marc Robin. Equilibrium Wage Dispersion with Worker and Employer Heterogeneity. *Econometrica*, 70(6):2295–2350, 2006. 1.1

Roy Radner and Andrew Schotter. The Sealed-Bid Mechanism: An Experimental Study. *Journal of Economic Theory*, 48:179–220, 1989. 1.1, 2.5

Alvin E. Roth and Xiaolin Xing. Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions. *American Economic Review*, 84(4):992–1044, 1994. 29

Emmanuel Saez and Gabriel Zucman. Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data. *Mimeo*, 2014. 1

Mark A. Satterthwaite and Steven R. Williams. Bilateral Trade with the Sealed Bid k-Double Auction: Existence and Efficiency. *Journal of Economic Theory*, 48:107–133, 1989. 1.1, 2.4, 2.5

Jae Song, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. Firming Up Inequality. *mimeo*, 2016. 1.1

Jean Tirole. From Bottom of the Barrel to Cream of the Crop: Sequential Screening with Positive Selection. *Econometrica*, Forthcoming, 2016. 2.3

Donald M. Topkis. *Supermodularity and Complementarity*. Princeton University Press, 1998. 2

Joel Watson. Perfect Bayesian Equilibrium: General Definitions and Illustrations. *mimeo*, 2016. 43

Martin L. Weitzman. The 'Ratchet Principle' and Performance Incentives. *Bell Journal of Economics*, 11(1):302–308, 1980. 2.3

Steven R. Williams. Efficient Performancy in Two Agent Bargaining. *Journal of Economic Theory*, 41:154–172, 1987. 3.3

9 Tables and Figures

Table 1: Summary Statistics, TaskRabbit

	Post Price (mean)	Priv. Auction (mean)	T-Stat (Post− Auct.)	P-Value (Post− Auct.)	Post Price (N)	Private Auction (N)
Initial wages ($)	41.1	61.0	-43.3	0.00	> 100k	> 100k
Share w/ raise	0.13	0.15	-23.7	0.00	> 100k	> 100k
Share incorp. bus.	0.10	0.11	-5.79	0.00	> 100k	> 100k
Emp. ever bonus	0.32	0.31	3.86	0.00	> 100k	> 100k
Share emp. (> $150k)	0.45	0.45	-0.38	0.71	> 100k	> 100k
Share w/ 1 posting	0.46	0.44	7.13	0.00	> 100k	> 100k

(a) Summary statistics of all jobs. Observation numbers are intentionally obscured at the request of TaskRabbit.

(mean)	Separated (mean)	Co-located (Sep. − Co.)	T-Stat (42% separate)	Accpt. Bids (44% separate)	Tasks (43% separate)	Employers
Initial wages ($)	49.8	58.9	-5.00	2787	1064	663
Distance from top bid (%)	0.43	0.33	2.37	2787	1064	663
Number hired	2.60	2.69	-0.94	2787	1064	663
Received bids (Gini)	0.19	0.20	0.25	2787	1064	663
Chosen bids (Gini)	0.09	0.08	1.82	2787	1064	663
Positive ratings (Gini)	0.34	0.33	0.42	2787	1064	663
Share w/ raise	0.04	0.19	-9.96	2787	1064	663
Raise (unconditional)	0.04	0.12	-3.64	2787	1064	663
Share incorp. bus.	0.56	0.62	-1.27	2787	1064	663
Emp. ever bonus	0.27	0.60	-8.44	2787	1064	663
Share w/ 1 posting	0.52	0.66	-3.34	2787	1064	663
Female employer	0.56	0.62	-1.28	2787	1064	663

(b) Summary statistics of multi-worker jobs.

	N	Mean	Stand. Dev.	25th Perc.	Median	75th Perc.
Share posted price	336	0.43	0.10	0.37	0.41	0.46
Market age (months)	336	16.6	12.2	6.1	15.3	26.1
Emp. to worker ratio	336	2.52	0.96	1.87	2.36	3.00
Job match rate	336	0.46	0.11	0.41	0.48	0.53
Price ($)	336	56.1	8.69	52.0	57.2	61.0

(c) City-month level summary statistics

Table 2: Summary Statistics, Experiment

		Negotiable Pay			Non-Negotiable Pay		
	Others' Pay	Not Transparent	Transparent	T-Statistic	Not Transparent	Transparent	T-Statistic
		(mean)	(mean)	(diff)	(mean)	(mean)	(diff)
Workers	Prior transcription expertise	0.18	0.18	0.06	0.18	0.19	-0.27
	Age	37	38	-0.37	35	36	-0.98
	Share female	0.56	0.61	-0.85	0.48	0.49	-0.23
	Share w/ at least some college	0.96	0.92	1.38	0.93	0.90	1.48
	N	*164*	*122*		*272*	*242*	
Managers	Prior management expertise	0.26	0.18	0.89	0.27	0.21	0.95
	Age	37	38	-0.68	36	35	0.91
	Share female	0.57	0.68	-1.04	0.52	0.48	0.54
	Share w/ at least some college	0.92	0.87	0.88	0.95	0.88	1.81
	N	*53*	*38*		*89*	*73*	

Note: All workers and managers are from the U.S. and recruited through Amazon's Mechanical Turk.

Table 3: Bonuses Among Co-located Workers

	(1)	(2)	(3)	(4)	(5)	(6)
Dep.Var.	Any Raise (Yes = 1)			Final Pay (% Above Bid)		
Distance from Top Bid (%)	0.00954	0.0125	0.00984	0.741***	0.778***	0.959***
	[0.0112]	[0.0106]	[0.0110]	[0.150]	[0.143]	[0.0952]
Number participants (log)		-0.0826***	0.00159		-0.377	-0.221
		[0.0247]	[0.00182]		[0.230]	[0.220]
Length of task description (log)		-0.0132	-0.0139		-0.0459	-0.159
		[0.0283]	[0.00946]		[0.148]	[0.175]
Frequency of task category (log)		-0.0467	-0.0220		-0.0650	0.0298
		[0.0349]	[0.0231]		[0.102]	[0.116]
Required personal equipment		0.0563	-0.0724***		0.561	0.491
		[0.109]	[0.0275]		[0.956]	[0.947]
Business incorporated		-0.132**	-0.0226		0.134	0.390*
		[0.0580]	[0.0294]		[0.175]	[0.212]
Business status unclassified		-0.102*	-0.0243		-0.0881	-0.0119
		[0.0594]	[0.0200]		[0.137]	[0.171]
Worker months experience		2.668	0.0940		2.734	-4.661
		[1.792]	[0.117]		[5.114]	[6.721]
Prior # closed offers		-0.00930	-0.123**		-0.0867*	-0.0775
		[0.00863]	[0.0618]		[0.0509]	[0.0638]
Female (yes = 1)		-0.0143	-0.0952		-0.244**	-0.267**
		[0.0208]	[0.0630]		[0.115]	[0.131]
Constant	0.184***	0.705**	0.561**	0.327***	1.654	0.664
	[0.0228]	[0.310]	[0.227]	[0.0708]	[1.251]	[1.398]
Category FE		✓	✓		✓	✓
>1 hour overlap			✓			✓
Observations	1872	1872	1568	351	351	267
Clusters	462	462	394	181	181	131
R^2	0.000547	0.0931	0.0968	0.439	0.501	0.591

Notes: Each model is estimated by OLS. Col. 1 through 3 are linear probability models. An observation is an accepted worker-bid for jobs with co-located workers. Top bidders are excluded to test SF1, which states distance from highest bid is not predictive of a bonus. The dependent variable equals one if the particular worker earns more than their agreed to bid, and 0 otherwise. Col. 4 through 6 is restricted to those workers who receive strictly more than their bid. The dependent variable is the size of the raise, as percent above bid. The primary explanatory variable is the distance the initial bid is from the maximum bid selected. Standard errors are clustered at the level of the employer.

Table 4: Dispersion in Final Wages and Worker Surplus: Experimental & Observational Evidence

	TaskRabbit				Experimental Evidence					
	(1)	(2)	(3)		(4)	(5)	(6)	(7)	(8)	(9)
Dep. Var.	Final Pay (Gini)				Final Pay (Gini)		Hourly Wage (Gini)		Worker Surplus (Gini)	
Transp.	-0.0463**	-0.0648***	-0.111***	Transp.	-0.0232***	-0.0215***	0.0282	0.0233	0.243**	0.311**
(Co-loc.)	[0.0222]	[0.0220]	[0.0286]	(Co-loc.)	[0.00707]	[0.00802]	[0.0363]	[0.0376]	[0.118]	[0.141]
Constant	0.126***	0.141***	0.217**	Constant	0.0232***	0.114	0.144***	0.0933	0.115**	-0.164
	[0.0219]	[0.0220]	[0.0898]		[0.00707]	[0.0935]	[0.0219]	[0.215]	[0.0454]	[0.449]
Category FE		✓	✓	Mgr. Char.		✓		✓		✓
Employer FE			✓							
Mean D.V.	0.0829	0.0829	0.0829		0.0147	0.0147	0.154	0.154	0.204	0.204
Obs.	1046	1046	1046	Obs.	74	74	74	74	74	74
R^2	0.0180	0.0971	0.810	R^2	0.0794	0.171	0.00831	0.122	0.0729	0.146

Notes: Each model is estimated by OLS. An observation is a multi-worker local task. Standard errors are clustered at the employer level. The dependent variable in Col. 1 - 3 is the dispersion of the final pay in a TaskRabbit multi-worker job measured by the Gini Coefficient. Col. 4 -9 present experimental evidence. The dependent variable Col. 4-5 is the dispersion in final pay, in Col. 6 and 7 dispersion in hourly wages, and in Col. 8 and 9 dispersion in worker surplus, which we measured by subtracting employees' minimum willingness to accept the job (elicited using an incentive compatible method similar to that of Becker et al. (1964)) from the price paid. Manager characteristics in the experimental setting include gender, age, age squared managerial experience, and years of formal education.

Table 5: Endogenous Selection of Transparent Pricing

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.	Choose Transparent Posted Price (yes =1)					
Low Income Employer	0.0622**	0.0641**	0.0456*	0.0448*	0.0452*	0.0579*
	[0.0269]	[0.0272]	[0.0270]	[0.0269]	[0.0270]	[0.0329]
Employer Age		-0.00147***	-0.000511	-0.000524	-0.000536	-0.00109
		[0.000557]	[0.000524]	[0.000527]	[0.000528]	[0.000672]
Empl. Gender (Fem = 1)		-0.0130	-0.00901	-0.00926	-0.00921	-0.0144
		[0.0148]	[0.0130]	[0.0129]	[0.0129]	[0.0163]
Age of Marketplace					-0.000895	-0.0000381
					[0.000592]	[0.000714]
Constant	0.323***	0.390***	0.221***	0.225***	0.256***	0.215***
	[0.0259]	[0.0359]	[0.0538]	[0.0573]	[0.0602]	[0.0720]
Category FE			✓	✓	✓	✓
Metro FE				✓	✓	✓
Exclude 1st time users						✓
Mean Dep. Var.	0.423	0.423	0.423	0.423	0.423	0.423
Observations	35230	35230	35230	35230	35230	25964
Clusters	9554	9554	9554	9554	9554	5326
R^2	0.000432	0.00177	0.100	0.101	0.102	0.112

Notes: All columns are linear probability models estimated by OLS. An observation is a job post on TaskRabbit. Standard errors are clustered at the level of the employer. The dependent variable is equal to 1 if the employer chose to post the job using a transparent (public) posted price, and 0 if the employer chooses to accept private bids. The primary explanatory variable, low income, is an indicator equal to 1 if the employer earns less than the median earning household in each city, approximately $150k annually.

Table 6: Effect of Transparency on Employment, by Value of Labor

	TaskRabbit				Experimental Evidence		
	(1)	(2)	(3)		(4)	(5)	(6)
Dep. Var.	Vacancy filled (yes = 1)			Dep. Var.	Hired worker (yes = 1)		
Low Inc. X Transp. Price	0.0203	0.0229*	0.0409**	Low Value X Transp. Tx.	0.339***	0.286***	0.267***
	[0.0139]	[0.0128]	[0.0200]		[0.0881]	[0.0922]	[0.0898]
Low Inc.	-0.0210*	-0.0172*	-0.0187	Low Value	-0.224***	-0.166**	-0.169**
	[0.0114]	[0.0101]	[0.0140]		[0.0759]	[0.0811]	[0.0803]
Transp. Price	0.170***	0.132***	0.143***	Transparent Tx.	-0.0886*	-0.0258	-0.025
	[0.0112]	[0.0103]	[0.0164]		[0.0510]	[0.0565]	[0.0610]
Empl. Char.		✓	✓	Mgr. Char.		✓	✓
Category FE		✓	✓	Worker Char.			✓
City FE, Month FE, Mkt. Age			✓				
Mean Dep. Var.	0.650	0.650	0.650	Mean Dep. Var.	0.867	0.867	0.867
Observations	35230	35230	35230	Observations	197	197	196
Clusters	9554	9554	9554	Clusters	63	63	63
R^2	0.0352	0.0664	0.0745	R^2	0.0734	0.117	0.16

Notes: Each model is a linear probability model estimated by OLS. In Col. 1, through 3, an observation is a job posting on TaskRabbit. The dependent variable is equal to 1 if the job posting is matched to a worker before it expires on TaskRabbit. Col. 4, though 6 are from our experimental data. An observation is worker. The dependent variable is equal to 1 is this worker is hired, and 0 otherwise. Standard errors in all columns are clustered at the level of the employer.

Table 7: Share of Jobs with Fully Transparent Posted Price

	(1)	(2)	(3)	(4)
Dep. Var.	Proportion of Jobs with Posted Price			
Market age (months)	0.0109***	0.00894*	0.0119***	0.0102**
	[0.0000198]	[0.00426]	[0.00103]	[0.00383]
City FE	✓	✓	✓	✓
Calendar months FE	✓	✓	✓	✓
Number of posts per month (thousands)			✓	✓
Share of jobs in each category		✓		✓
Observations	417	417	417	417
Clusters	19	19	19	19
R^2	0.668	0.731	0.717	0.734

$^{*}p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Each model is estimated by OLS. An observation is a city-month in TaskRabbit. The dependent variable is the proportion of tasks that use the transparent posted price scheme. Standard errors are clustered at the city level.

Table 8: Gender gap in likelihood of raise rises in co-located, partially transparent jobs

	(1)	(2)	(3)	(4)	(5)	(6)
Dep.Var.	Any Raise (Yes = 1)			Final Pay (% Above Bid)		
Together	0.183***	0.164***	0.162***	0.159	0.113	0.216
	[0.0262]	[0.0232]	[0.0239]	[0.163]	[0.247]	[0.366]
X Female (yes = 1)	-0.0905***	-0.0670***	-0.0728***	0.155	0.314	0.367
	[0.0263]	[0.0239]	[0.0258]	[0.220]	[0.297]	[0.436]
Female (yes = 1)	0.00859	0.0423***	0.0430***	-0.0599	-0.191	-0.206
	[0.0153]	[0.0164]	[0.0152]	[0.158]	[0.220]	[0.349]
Exp. on platform (months)		0.00192	0.000912		-0.000453	-0.00261
		[0.00131]	[0.00130]		[0.00547]	[0.00700]
Prior # closed offers		-0.00848	-0.0104		-0.0777	-0.0662
		[0.00654]	[0.00697]		[0.0548]	[0.0718]
No. workers (log)		-0.0780***	-0.0646***		-0.288	-0.352
		[0.0190]	[0.0211]		[0.282]	[0.336]
Length of job post (log)		-0.0100	-0.0145		0.0169	0.0300
		[0.0218]	[0.0226]		[0.198]	[0.254]
Freq. of task category (log)		-0.0488	-0.0266		-0.219	-0.159
		[0.0336]	[0.0207]		[0.153]	[0.145]
Required equipment		0.0405	0.0683		0.595	0.501
		[0.0926]	[0.0981]		[0.894]	[0.901]
Business incorporated		-0.100**	-0.0911*		0.116	0.168
		[0.0459]	[0.0478]		[0.198]	[0.264]
Business status unclassified		-0.0885*	-0.0846*		-0.190	-0.202
		[0.0458]	[0.0478]		[0.179]	[0.240]
Constant	0.0408***	0.518*	0.345*	0.415***	2.839*	2.046
	[0.0117]	[0.289]	[0.206]	[0.142]	[1.632]	[1.617]
Category FE		✓	✓		✓	✓
> 1 hour overlap			✓			✓
Observations	2644	2644	385	385	288	288
Clusters	605	605	195	195	139	139
R^2	0.0419	0.104	0.110	0.00405	0.0617	0.0522

Notes: Each model is estimated by OLS. Standard errors are clustered at the level of the job. An observation is a worker-bid in a co-located, multi-worker job on TaskRabbit using a private auction. Col. 1 through 3 are linear probability models. The dependent variable equals 1 if the particular worker earns more than their initial bid, and 0 otherwise. Col. 4 though 6 are restricted to those workers who receive more than their bid. The dependent variable is the size of the raise, as percent above bid. The primary explanatory variable is an indicator for whether the job entails co-location of workers (a measure of partial transparency) and its interaction with the gender of the worker.

Table 9: Worker-Bid Level Pay Compression

		(1)	(2)	(3)	(4)
	Dep. Var.	Final Pay (% over bid)			
	Distance from Top Bid (%)	0.384***	0.385***	0.388***	0.283**
		[0.137]	[0.137]	[0.138]	[0.135]
	X Separated	-0.364***	-0.366***	-0.367***	-0.275**
		[0.139]	[0.140]	[0.140]	[0.137]
	X Virtual	-0.306**	-0.307**	-0.304**	-0.189
		[0.141]	[0.141]	[0.142]	[0.141]
	Separated	-0.00228	0.00307	0.0235	-0.0168
		[0.0512]	[0.0520]	[0.0420]	[0.0565]
	Virtual	-0.0359	-0.0289	0.00416	-0.112
		[0.0493]	[0.0498]	[0.0484]	[0.0765]
Worker Traits	Worker months experience		0.849	0.584	-0.204
			[1.439]	[1.304]	[1.252]
	Prior # closed offers		-0.000425	-0.00305	0.00822
			[0.00751]	[0.00839]	[0.00884]
	Female (yes = 1)		-0.0288	-0.00725	-0.0125
			[0.0200]	[0.0198]	[0.0216]
Job& Empl. Traits	No. workers (log)			-0.0501**	-0.0575**
				[0.0201]	[0.0240]
	Length of job post (log)			0.0424	0.00189
				[0.0390]	[0.0197]
	Freq. of task category (log)			-0.0389	0.00654
				[0.0346]	[0.0194]
	Required equipment			0.383	0.0210
				[0.493]	[0.0997]
	Business (incorp.= 1)			-0.0495	
				[0.0659]	
	Business (unclassified)			0.0654	
				[0.0845]	
	Constant	0.0616	0.0698	0.231	0.173
		[0.0461]	[0.0526]	[0.336]	[0.233]
	Category FE			✓	✓
	Employer FE				✓
	Mean dep. var. (all multi-worker jobs)	0.115	0.115	0.115	0.115
	Observations	3235	3235	3232	3232
	Clusters	464	464	464	464
	R^2	0.172	0.173	0.198	0.583

Notes: Each model is estimated by OLS. An observation is an accepted worker-bid for a multi-worker job on TaskRabbit. Standard errors are clustered at the level of the employer. The dependent variable is the size of the raise, as percent above the worker's initial bid. The primary explanatory variable is an interaction between the distance the initial bid is from the maximum bid selected and a categorical variable for co-located, separated or virtual workers.

Table 10: Lower Bids When Job Post Mentions Price

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. Var.		Offers (log)			Final Pay (log)	
Price in Job Post	-0.224***	-0.0953***	-0.0954***	-0.335***	-0.104***	-0.134***
	[0.0155]	[0.0144]	[0.0136]	[0.0167]	[0.0158]	[0.0144]
Exp. on platform (days)		0.304***	-0.407***		0.170***	-0.127***
		[0.0143]	[0.0217]		[0.0280]	[0.0456]
Mean rating in category		0.0298***	0.0420***		0.0842***	0.0795***
		[0.00748]	[0.00740]		[0.0198]	[0.0201]
# reviews in category		-0.355***	-1.275***		1.004***	-0.713***
		[0.0885]	[0.0863]		[0.175]	[0.172]
Mean rating overall		0.0656***	-0.00792		0.0728***	-0.00430
		[0.00818]	[0.0115]		[0.0233]	[0.0332]
# ratings overall		0.755***	1.797***		0.573***	1.756***
		[0.0393]	[0.0557]		[0.0783]	[0.123]
Female (yes = 1)		-0.117***			-0.0743***	
		[0.00249]			[0.00533]	
Exp. on platform (sqr.)		-0.266***	0.0237		-0.219***	-0.123***
		[0.0137]	[0.0195]		[0.0247]	[0.0412]
Mean rating in category (sqr.)		-0.00423***	-0.00872***		-0.0137***	-0.0157***
		[0.00150]	[0.00148]		[0.00399]	[0.00404]
# reviews in category (sqr.)		-1.386***	0.552**		-3.783***	-0.848*
		[0.224]	[0.217]		[0.495]	[0.448]
Mean rating overall (sqr.)		-0.0189***	-0.00167		-0.0147***	0.00223
		[0.00164]	[0.00231]		[0.00468]	[0.00668]
# ratings overall (sqr.)		-0.371***	-0.983***		-0.198*	-0.827***
		[0.0564]	[0.0719]		[0.105]	[0.154]
Constant	3.770***	3.387***	3.838***	3.612***	3.104***	3.526***
	[0.00299]	[0.173]	[0.00453]	[0.00303]	[0.274]	[0.00926]
Category FE		✓	✓		✓	✓
Worker FE			✓			✓
Observations	>100k	>100k	>100k	>100k	>100k	>100k
Clusters	>100k	>100k	>100k	>100k	>100k	>100k
R^2	0.00294	0.0831	0.277	0.00875	0.159	0.369

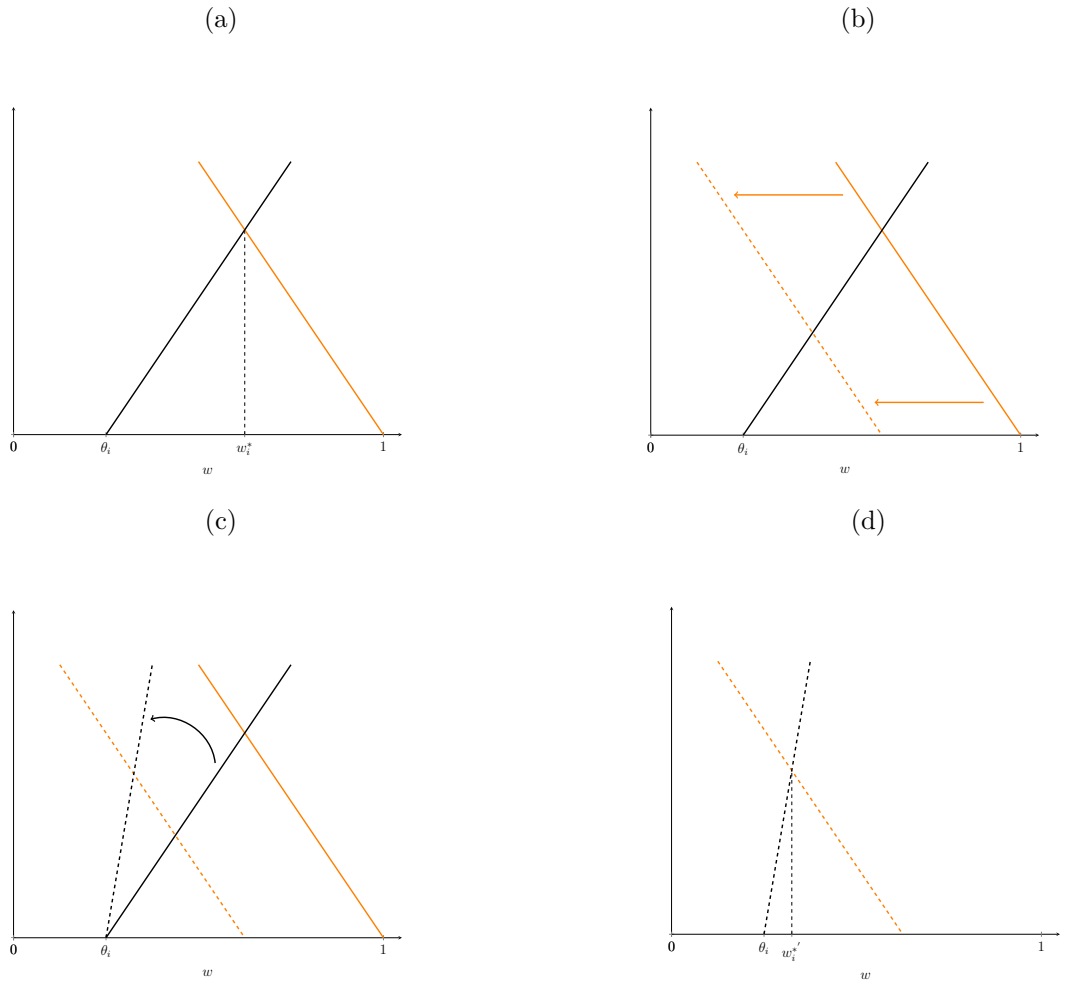
Notes: Each model is estimated by OLS. An observation is a worker-bid on TaskRabbit. The dependent variable in Col. 1 through 3 is the log bid received, and in Col. 4 through 6 is the log final pay among those bids that were accepted. The "price in job post" indicator variable refers to any mention of the price or price range embedded in the job description. Only job posts that accept private bids are included in these regressions. Standard errors are clustered by work item.

Table 11: Higher Profits Under Transparency

		(1)	(2)	(3)
	Dep. Var.	Profits ($)	Profits ($)	Profits (log)
	Full Transparency	1.970**	2.004***	0.536***
		[0.754]	[0.755]	[0.181]
Manager Traits	Female		-0.932	-0.338**
			[0.621]	[0.162]
	Age		0.287	0.0566
			[0.204]	[0.0571]
	Age Sqr.		-0.00299	-0.000509
			[0.00255]	[0.000710]
	Some college		-2.009	-0.422
			[1.570]	[0.295]
	Mgr. exp. (interm.)		0.643	0.199
			[0.839]	[0.219]
	Mgr. exp. (adv.)		1.193	0.428*
			[0.865]	[0.231]
	Transcr. exp. (interm.)		-2.242	-0.344
			[1.530]	[0.331]
	Transcr. exp. (adv.)		-2.216	-0.548
			[1.867]	[0.387]
	Constant	2.777***	0.256	0.192
		[0.708]	[3.802]	[1.153]
	Observations	90	90	90
	R^2	0.0901	0.229	0.234

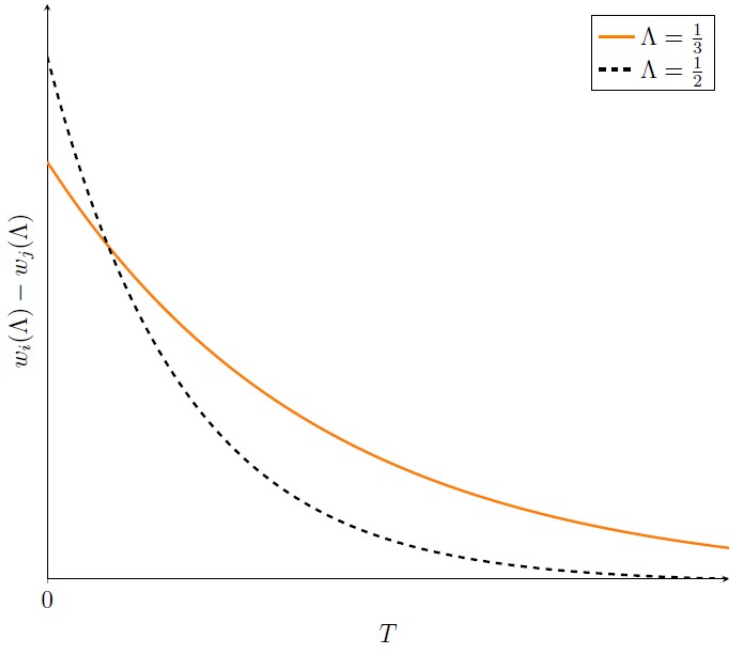
Notes: Each model is estimated using OLS. An observation is a manager in our experiment. The dependent variable in Col. 1 and 2 is the total profit a manager earns. Col. 3 is the log profit, including the base payment for participating. Robust standard errors are in square brackets.

Figure 3: Effects of increasing Λ on worker and firm strategies



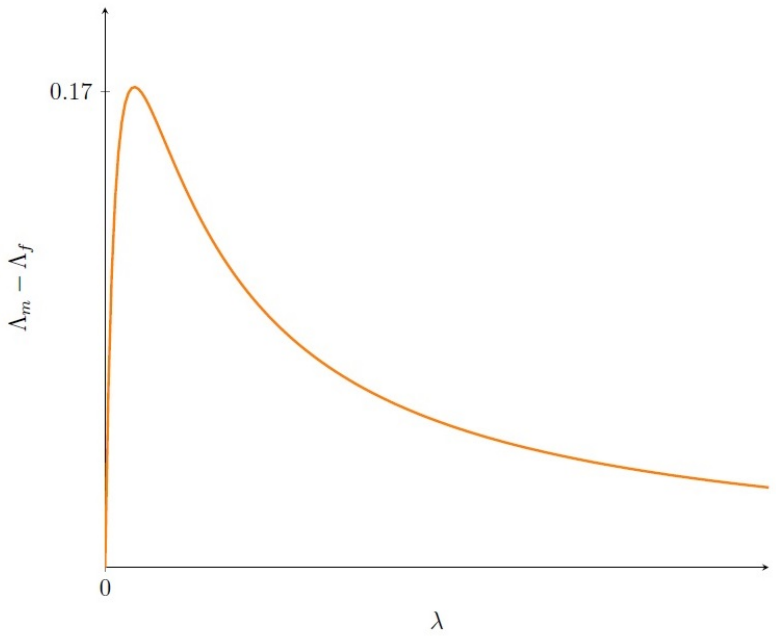
Notes: (a) By Equation 5, worker i picks an initial wage offer w_i^* that equalizes $\frac{w-\theta_i}{1-\Lambda}$ (the black, upward sloping line) and $\frac{1-\bar{F}(w)}{\bar{f}(w)}$ (the orange, downward sloping line). (b) The **demand effect** of increasing Λ from 0 to $\frac{3}{4}$ reduces $\bar{w}$ for each v, shifting $\frac{1-\bar{F}(w)}{\bar{f}(w)}$ to the left. (c) The **supply effect** of increasing Λ from 0 to $\frac{3}{4}$ increases the slope of the function $\frac{w-\theta_i}{1-\Lambda}$. (d) The **supply** and **income effects** combine to reduce the initial wage offer of worker i to $w_i^{*\prime}$ when Λ increases.

Figure 4: Expected difference in equilibrium wages T periods after entering the market



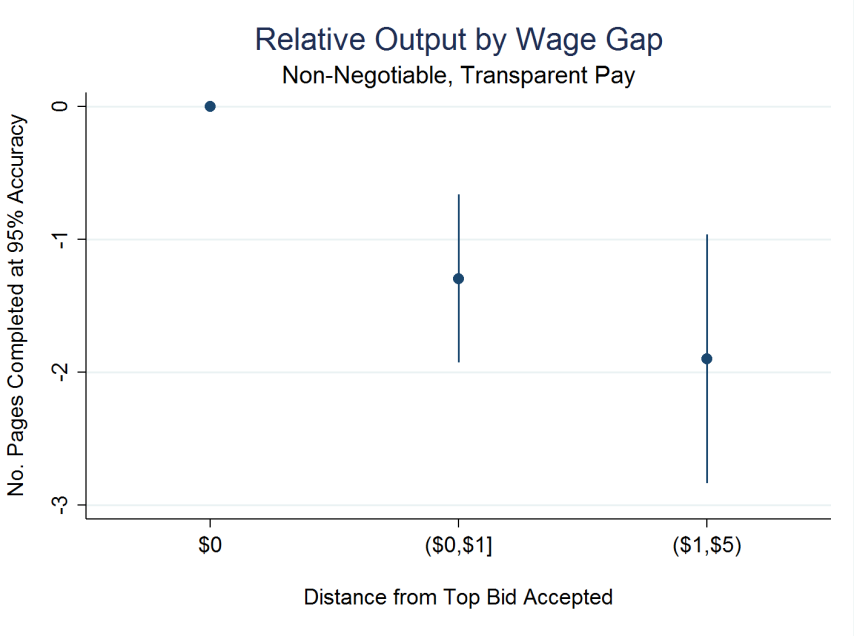
Notes: Figure 4 shows the expected difference in the wage of two workers, i and j T periods after each has entered the market when $\theta_i > \theta_j$. The dashed (black) curve represents this difference when $\Lambda = \frac{1}{2}$, $\rho + \delta = 1$, $r = s = 1$, and the solid (orange) curve represents this difference when $\Lambda = \frac{1}{3}$, $\rho + \delta = 1$, $r = s = 1$. Although the dashed curve is initially above the solid one, the two curves satisfy a single-crossing condition in t.

Figure 5: Difference in de facto arrival rate of wage info between genders as a function of λ



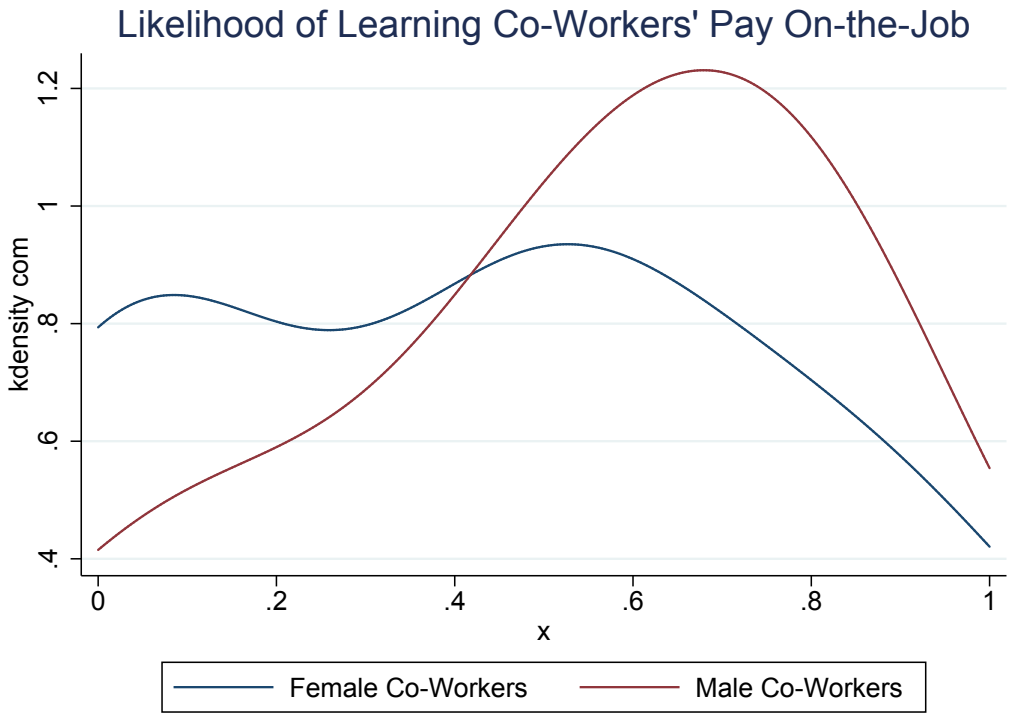
Notes: Figure 5 plots λ (horizontal axis) and the difference in the de facto rate of information arrival between men and women. This difference is initially increasing in λ, but after a single peak, it decreases toward 0. Parameters used: $q = .5$, $\alpha_m = 4$, $\alpha_f = 2$, $\rho + \delta = 1$.

Figure 6: Productivity consequences of transparency when pay is non-negotiable



Notes: This graph displays the coefficients estimated by regressing the number of pages completed on a flexible functional form for the gap in pay, using OLS. Above we grouped the distance from highest bid accepted into three bins, exactly equal, between 0 and $1 in distance, and $1 upwards. The data include 150 managers, and 267 workers who bid less than or equal to the $5 budget.

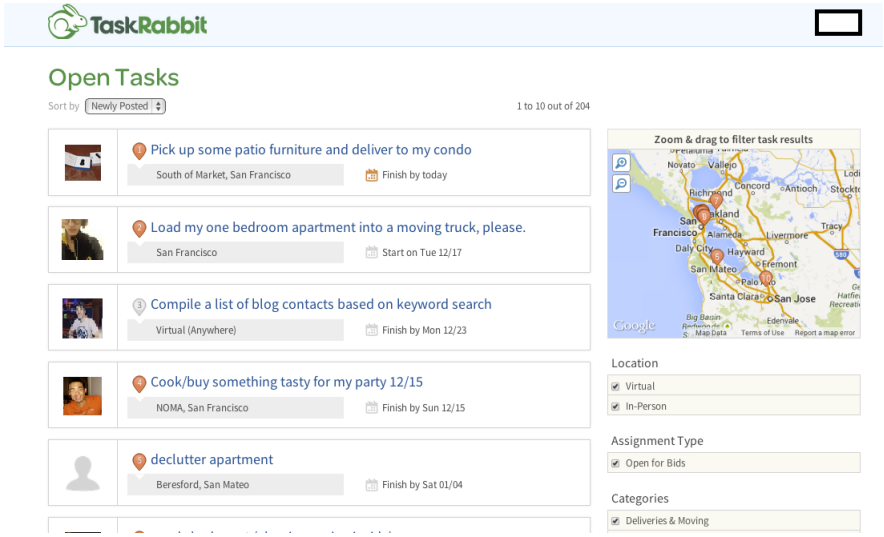
Figure 7: Expectations of learning co-worker pay on-the-job



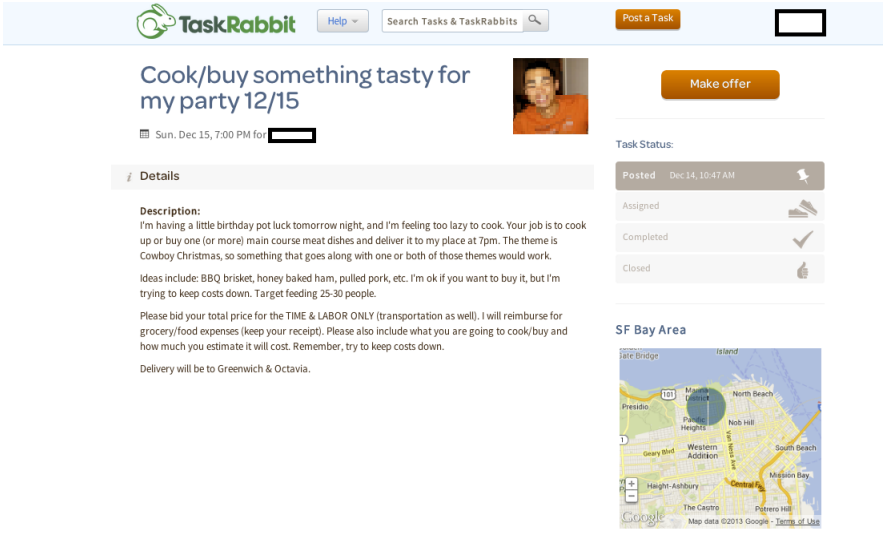
Notes: This figure is a kernel density constructed from 5,000 responses from online workers who read through job descriptions on TaskRabbit and answered questions about the likelihood that two co-workers would compare notes about their pay after meeting for the first time on-the-job. We experimentally changed the names of the co-workers to signal the co-workers were either male or female. Here we plot female respondents reporting about female co-workers, and male respondents about male co-workers.

Figure 8: TaskRabbit Online Interface for Workers

(a)

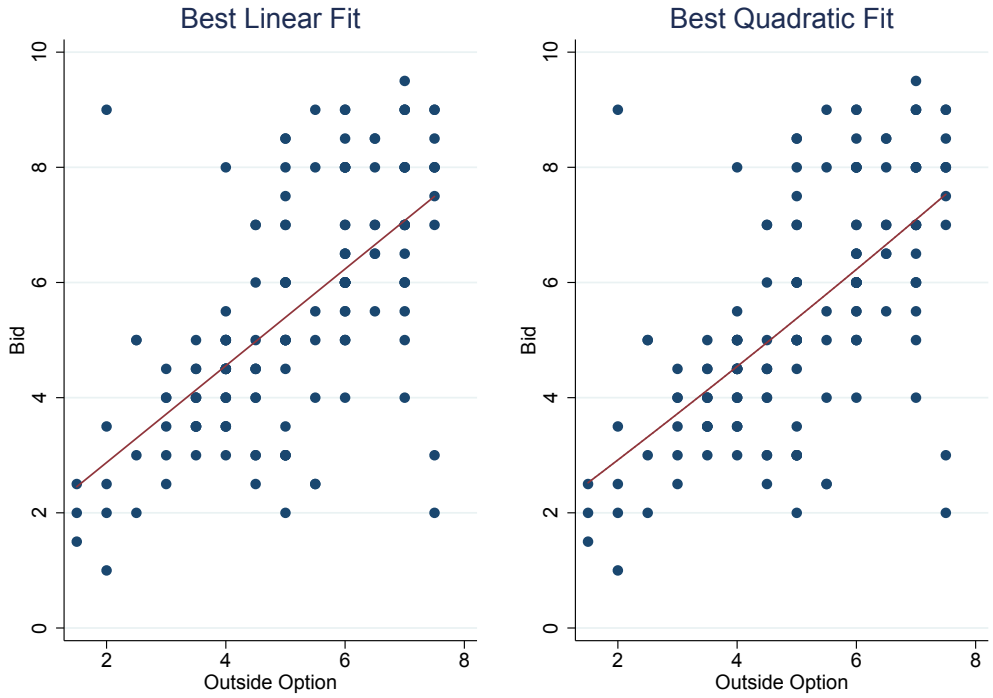


(b)



Notes: Panel (a) displays a list of job postings that a worker can see. Panel (b) gives the details posted by the employer about one of the jobs from the job listings page. Screenshots taken on December 14th, 2013. Faces and identifiable information have been intentionally blurred. A similar figure appears in Cullen and Farronato (2016).

Figure 9: Bids as function of Willingness to Accept



Notes: Each panel plots the outside option of a participant (horizontal axis) as measured by our BDM procedure against the participant's bid on the job for completion of a page of transcription (vertical axis), both at a minimum accuracy of 95%. In the first panel, we fit the data to a best linear fit of outside option, and in the second, we regress on both outside option and outside option squared. The best fit curves are nearly identical.

Appendix

A Tests of alternative explanations

In this section we assess mechanisms other than communication about pay per se that could explain the distinct wage setting behavior we observe when workers are co-located. Chief among them are productivity spillovers, either observed or perceived. Under a pay-for-performance framework, an employer may assign more compressed wages to workers if their performance converges or if the employer cannot attribute the output to individual workers.

Perceived productivity differences, as the explanation for the wage compression we observe, requires that (1) employers compensate workers according to their on-the-job assessed performance and (2) assessed performance of co-workers is less dispersed when workers are together.

We find evidence that a component of pay that reflects on-the-job performance using a measure of performance constructed from back-end administrative data (effective percent positive score (EPP), detailed in Nosko and Tadelis (2015)). However, we do not find empirical evidence to support (2). Our performance measure is no more dispersed or compressed when workers are co-located. In Table 14, the dependent variable is the dispersion of ratings given to workers at the conclusion of the job, expressed as the Gini coefficient. An indicator variable of whether these workers operate together or separately proves uninformative about the final dispersion of ex-post ratings. Since we find no evidence that employer evaluations (or ex-post ratings) converge among workers when they are together, it is unlikely that productivity drives the wage compression we observe.

More generally, there is a weak relationship between bids and productivity both theoretically and empirically. When we can observe productivity directly in our field experiment we find a small and insignificant relationship between productivity and bids. A high productivity type might have both lower costs of effort and higher opportunity costs. While our measures of productivity on TaskRabbit are strong predictors of real outcomes (return customers) they do not explain much of the variance in market wage, so any systematic pattern of spillovers does not necessarily raise the performance of the low bidder or the pay of the low bidder per se. In other words, a model of positive spillovers where the most productive worker pulls up the performance of the least productive worker, would not imply that the lowest bidder improves performance per se.

Table 12: EPP predicts other productivity measures

	(1)	(2)
Dep. Var.	Employer returns	Positive rating
EPP	1.591***	5.858***
	[6.154]	[20.19]
Ex-Ante mean rating	0.955**	0.876***
	[-2.456]	[-5.611]
Prior # closed offers	1.075***	0.857***
	[6.562]	[-11.72]
Constant	0.168	1.035
	[-1.156]	[0.0261]
Category FE	✓	✓
Worker characteristics	✓	✓
Job Characteristics	✓	✓
Observations	> 100k	> 100k

Exponentiated coefficients; t statistics in brackets

Notes: Each model is estimated using maximum likelihood assuming extreme value type-1 distributed errors (logistic regression). An observation is a matched worker-job in TaskRabbit. In Col. 1 the dependent variable equals 1 if the employer returns to the platform after the job is completed, giving her the option to rate the worker. The dependent variable in Col. 2 is equal to 1 if the worker receives a positive review after the job is complete, 0 otherwise. Positive review is defined as either a 4 or 5 on the 5 star scale. Standard errors are clustered at the job level. T-statistics are reported in brackets beneath the point estimate. Job characteristic controls include category fixed effects and proxies for transparency of the job requirements, including the length of description and frequency of posts in same category. We also include the number of bidders (log) and equipment requirements.

Table 13: Productivity measure (EPP) predicts ex-post pay but not ex-ante pay

	(1)	(2)	(3)
Dep Var:	Bid (log)	Raise (%)	Raise (%)
Ex-ante EPP	0.00960	0.0771*	0.229**
	[0.0461]	[0.0431]	[0.0927]
× Separate places			-0.0904
			[0.0821]
× Virtual			-0.162**
			[0.0797]
× Single Worker			-0.149*
			[0.0905]
Entry Month FE	✓	✓	✓
Category FE	✓	✓	✓
Worker characteristics	✓	✓	✓
Job Characteristics	✓	✓	✓
Mean Dep. Var.	3.63	0.129	0.129
Observations	>100k	>100k	>100k
R^2	0.238	0.00583	0.00456

Notes: All models are estimated by OLS. An observation is the bid from a worker assigned to a job on TaskRabbit. The dependent variable is the log bid in Col. 1 and ex-post pay out above and beyond the initial bid in Col. 2 and 3. Standard errors are closed at the level of the worker.

Table 14: Dispersion in Perceived Worker Performance

	(1)	(2)	(3)
Dep. Var.	Worker Performance Ratings (Gini)		
Transp. (Co-loc.)	-0.00946	-0.00788	0.0294
	[0.0235]	[0.0204]	[0.0542]
Constant	0.406***	0.675***	0.406***
	[0.0201]	[0.0637]	[0.0921]
Category FE		✓	✓
Employer FE			✓
Mean Dep. Var.	0.435	0.435	0.435
Observations	1064	1064	1064
Clusters	618	618	618
R^2	0.000260	0.0796	0.725

Notes: Each model is estimated by OLS. An observation is a multi-worker job in TaskRabbit. Standard errors are clustered at the employer level. The dependent variable is the dispersion in ratings received after work is completed, measured by a Gini coefficient.

A.1 Importance of a communication channel for wage-setting

Our administrative data suggests that communication about relative pay drives wage compression. We reaffirm this by replacing our binary measure of communication potential (co-location) with a continuous measure of the likelihood of learning about co-worker pay. We elicit expectations about communication and other activities on the job in a survey of similar workers. Survey participants report the likelihood that co-workers discuss pay is 47 percent on average with substantial heterogeneity among co-located jobs. The likelihood of learning about the pay of co-workers is a uniquely strong predictor of wage compression.

We replicate our main empirical specification, Eq.9, and replace our indicator for co-located jobs with the continuous measure of communication potential. We also add to our specification estimates of the likelihood that co-workers will influence each others' output and the odds that the employer will be unable to observe individual output. We predict the impact of communication on final pay is a function of how far apart bids are.[40] The impact of spillovers and the ability to monitor are a function of relative productivity. In our estimation of the importance of these job attributes on ex-post pay, we include a two-way interaction between communication potential and difference in bids, and we include a three-way interaction between potential for influence, inability to observe individual output, and

[40]This does not allow for relative pay concerns to be conditional on perceived quality differences. To the extent that differences in productivity relax fairness concerns, it will attenuate our estimates (which are sizable).

the productivity gap between workers.[41]

In Table 15, we find that increasing the likelihood that two workers on a particular job will communicate pay has a positive effect on the eventual pay of the low bidders. When the communication of pay is certain, a 1 percent increase in the ex-ante pay gap translates into a 1 percent increase in ex-post pay for the low bidder. When the chance of learning pay is expected to be fifty percent (as is true of the average job), a 1 percent increase in ex-ante pay difference results in a half of a percent increase in ex-post pay (nearly equivalent to our estimate of the average effect when co-located).

The potential role for positive spillovers[42] raises the average ex-post pay but does not increase wage compression. Nor does the ability to monitor affect wage compression. The underlying reason we see negligible effects of spillovers and observability on relative pay likely has more to do with the disconnect between pay and performance in this market.

A.1.1 Vignettes

Below we describe a brief posting for a real job that takes place in a real city. The employer asks workers on an online labor platform to submit the price they want to complete the job.

Image that two people, Sam and Alex, offer two different prices to do the job. Sam and Alex haven't met before and they don't know the other's price for the job because they submitted these offers from their own computers. The employer chooses Sam and Alex on the platform, and they arrange a time to do the work.

Below is a description of the job. Please read the description and answer the questions about what happens when Sam and Alex begin working.

[Insert job description]

Answer on a scale of 0 through 10. A value of 0 means they definitely will not. A value of 1 means the odds are 1 in 10. In other words, if it happened 10 times, they would probably only learn about each other's pay on one of those occasions. A value of 10 means that they would learn about it every time.

1. How likely do you think Sam and Alex will discover what the other one is being paid for the job?

2. Is this the type of job where either Sam or Alex could be a leader or role model and

[41] Measured as the probability of any influence or perfect visibility of individual output and the predicted market value of productivity attributes, respectively.

[42] The likelihood of influencing a co-worker according to survey respondents.

encourage the other to do a better or worse job?

3. Will the employer be able to see what each of them individually contributed?

4. What is the likelihood that either Alex or Sam would be hired by the same employer in the future for something similar?

5. How long would this work take each person?

6. How long would they we working together, if at all?

Table 15: Linking Crowd-Sourced Job Attributes to Renegotiation

	(1)	(2)
Dep Var	Δy_{ijt} : $ex-post$ pay relative to initial bid	
		Main Effect
Distance from Top Bid (%)	-0.162***	-0.160***
	[0.0536]	[0.0537]
Distance from Most Productive	-4.521*	-0.204
	[2.522]	[0.152]
Communicating pay	-0.250	-0.135
	[0.207]	[0.311]
Unobservable indiv. effort	0.0538	0.979
	[0.311]	[1.474]
Role for leadership	0.603*	-0.377
	[0.345]	[1.066]
Role for leadership × Unobservable indiv. effort	-0.225	-2.140
	[0.846]	[3.086]
		Differential Effect
Distance from Top Bid (%)	1.070***	1.064***
(× Communicating pay)	[0.164]	[0.166]
Distance from Most Productive	8.974	0.294
(× Peer spillovers)	[7.327]	[0.395]
Distance from Most Productive	12.54	0.716
(× Unobservable indiv. effort)	[8.372]	[0.559]
Distance from Most Productive	-4.547	-1.106
(× Role for leadership × Unobservable indiv. effort)	[14.47]	[1.288]
Job characteristics	✓	✓
Worker characteristics	✓	✓
Mean dep. var.	0.143	0.143
Observations	886	886
R^2	0.851	0.852

Each model is estimated by OLS. An observation is a successful worker bid. The dependent variable is the difference between final pay and initial bid expressed as a percentage raise. The main explanatory variable "likelihood of learning co-worker pay" is derived from a large survey of jobs. We calculate the mean score of survey participants (25 for each job description). Column (2) replaces the predicted wage of productivity types with a raw measure of life time effective percent positive rating (EPP). The average likelihood of communicating pay in this sample is 0.47; observability of effort 0.67 and mean potential for peer spillovers 0.41. Distance from the most productive is on average 0.25 (EPP) and 0.02 (productivity component of wage) with skew greater than 2. We re-sample from survey responses without replacement to bootstrap standard errors. Worker characteristic controls include months since entering the platform, number of prior jobs, and prior mean rating. Job characteristic controls include category fixed effects and proxies for transparency of the job requirements, including the length of description and frequency of posts in same category. We also include the number of bidders (log).

A.2 Strategic bidding, worker selection, and unanticipated transparency

For a causal interpretation of the effect of co-location on ex-post relative wages in our TaskRabbit population, we must show the composition of workers is similar across settings as are worker' bids. Prima facie evidence supports these assumptions. Multi-worker tasks comprise fewer than 5 percent of posted jobs and workers are often unaware that more than one vacancy exists even when it does. Additionally, employers rarely have more offers that

the number necessary to complete a multi-worker job. Here we offer more empirical tests.

We observe that the mean and dispersion of bids received are similar across job settings. We also find that dispersion in selected offers is no different across setting. Irrespective of work setting, employers select bids that exhibit roughly one-third of the dispersion of offers received.

As another test of our assumptions that workers, in this particular environment, do not bid strategically in anticipation of learning pay, we split a sample of co-located jobs by whether or not the employer explicitly mentions that the tasks require multiple people (eg. "we need two people to load boxes" vs "load boxes between 12-2p") and find almost all worker characteristics are not statistically different. A mention of more than one worker does not result in an increase the bids, even among those that are more experienced or have higher mean ratings.

To what extent is transparency expected? Posted price jobs very clearly offer transparency compared to bid acceptance jobs. But this does not appear to be the case for co-located jobs compared to separated or single worker jobs. We observe, as econometricians, the number of workers hired for each task and whether they are co-located or not, but this information is frequently not available to users on the platform; 54% of job postings for co-located, multi-worker jobs *do not* reveal to workers that there are other workers on the job. In these jobs, workers are unlikely to be able to anticipate transparency. As we observe significant wage compression caused by worker co-location, employers do not appear to anticipate transparency either. If they did, our framework describes a strong incentive to always inform workers of co-location. If workers expect to receive wage information of their peers, they will lower their bids, which benefits the employer; unanticipated transparency is strictly more harmful to the employer than anticipated transparency. Even among the jobs that specifically make mention of multiple co-located workers, transparency seems unanticipated by both employers and workers. As only 5% of jobs are multi-worker, users (especially employers who are short lived compared to workers) may not have enough experience with these jobs to anticipate the higher level of transparency and the effect of this transparency on wage bargaining.

A.3 Market Unraveling

We find evidence that TaskRabbit markets unravel toward the use of posted price by more employers in Section 4.8, which supports the finding of Corollary 2.

We discuss possible alternative explanations for this market trend toward posted price, and why we do not believe these to be plausible explanations of our observations.

One alternative explanation for this trend is that employers initially accept bids to learn about workers' outside options and in subsequent tasks use a posted price. We do not believe

this to be a convincing explanation for this observation because employers are short-lived in TaskRabbit. The majority of employers only post a single task on the platform, and the vast majority of employers post no more than three tasks. Nearly 80% of employers who participate in the platform *do not* experiment, that is, they use either posted price or bid acceptance for all of their tasks. Finally, even persisting employers are relatively short-lived; no employer is active in the platform for more than one year. Given the four year time horizon of our data, it is likely that the composition of employers within early-adopting cities to have reached steady state. The pattern of a linear move toward posted prices, therefore, seems unlikely to be due to experimentation.

Another alternative is put forth by Einav et al. (2013). They find that eBay's auction format became much less used than its posted price format between 2003 and 2009. They argue that this is primarily driven by a change in user preferences. In 2003 there were not many exciting internet alternatives, and so buyers preferred the fun associated with bidding in auctions. But by 2009 with the advent of Web 2.0 websites like youtube.com and facebook.com, there were better avenues for entertainment on the internet. Could a similar phenomenon be occurring in TaskRabbit? Again, we do not believe so. Our data sample (albeit for a different service) begins around the time that the sample of Einav et al. (2013) ends, certainly after the popularization of Web 2.0 and plenty of entertainment websites. Second, our time horizon is relatively short compared to theirs, and we observe a large move toward posted prices. Only a drastic change in preferences over a short period of time could explain this. Finally, and most convincingly, TaskRabbit had a staggered entry strategy into different markets, and therefore, we observe wide variance in market age. Despite this, we observe a strong linear trend toward posted price in markets of different ages. This is on display in Figure 2. Although changing preferences cannot be completed ruled out in TaskRabbit data, a mechanism such as Einav et al.'s does not seem likely to lead to the move toward posted price in TaskRabbit.

B Theoretical Appendix

B.1 Firm acceptance and rejection of offers

We introduced the game as one in which the firm selects a single $\bar{w}$ and is bound to that for all time. More realistically, the firm may be able to accept offers on a case-by-case basis. In this section, we show that generalizing the game and restricting our attention to a class of time consistent equilibria does not change the analysis.

Ammendments to the timing of the stage game are straightforward. Instead of selecting $\bar{w}$ at $t = 0$, the firm selects "accept" or "reject" for each offer as it receives it. By accepting,

the firm is locked in to paying the agreed upon wage until the worker departs or makes another offer, and if the firm rejects, then the worker is inelligible to work at the firm.

As we are interested in the effect of transparency on wage negotiation, learning about the wages of others must convey information about the wage a worker can request. Intuitively, we want to use an equilibrium refinement like Markov perfection, as this includes subgame perfection (so that the firm cannot make non-credible threats of refusing to accept certain wage offers) and time consistency (seeing the wage of a higher paid co-worker means that a worker knows she can receive that wage if she offers it to the firm). Unfortunately, Markov perfect equilibria are not well-defined in our setting.[43] Formally, we study equilibria satisfying **A0, A1-3**, **A4'**, and **A5**. We define **A0** and **A4'** below.

A0 The firm selects some function $\bar{w}(v)$, and accepts all offers $w_{i,t} \leq \bar{w}$ for any worker i and any time t, and rejects all others.

A4' Let $w_t^{\sup}$ be the highest wage paid by the firm at time t if the worker observes wages at time t, and 0 otherwise. Off path, each worker i believes with probability 1 that the firm will accept any offer she makes that is no more than $\max\{w_i^*, w_t^{\sup}\}$ and will reject all greater offers.

A0 restricts attention to firm strategies that set a maximum wage $\bar{w}$ that is constant *across* workers and over time *within* worker. This assumption that the firm's strategy is time-consistent within worker is a Markovian restriction; a firm can condition its acceptance strategy on v, previous offers made by the worker, and the history of the game. Note however, that given the constant inflow and outflow of workers, the only payoff relevant factor determining the state of the game from the firm's point of view is v. Furthermore, this Markovian assumption is necessary to understand the effects of pay transparency and worker bargaining. Because each worker is infinitesimally small, without any restriction, the firm could essentially negate pay transparency by refusing to renegotiate with workers. For example, the firm could play a strategy that defines some $\bar{w}_{i,t}(v)$, which is the maximum wage it will accept from each i at time t. The firm could set $\bar{w}_{i,t} = v$ and $\bar{w}_{i,t'} = w_i^*$ for all $t' > t$, which corresponds to the "full secrecy" world of $\lambda = 0$ we present later. Without this restriction, it is also possible to construct "sun spot" equilibria in which $\bar{w}_t(v)$ is a step function in t, that is at some time t' the firm's maximum willingness to pay jumps upward.

The restriction that the maximum accepted offer is equal across workers is motivated by the assumption that the firm cannot wage discriminate against workers as it does not observe outside options.[44] As we have limited our study to equilibria in which the firm's

[43] Watson (2016) discusses some issues of equilibrium refinement in games with infinite action spaces.

[44] Massachusetts recently passed a law prohibiting firms from asking potential employees their current salaries during job interviews (http://www.nytimes.com/2016/08/03/business/dealbook/wage-gap-

willingness to pay is constant over time within worker, if the firm had a different willingness to pay across workers, it would imply that the firm has a different willingness to pay for two workers i and j at the moment each of these workers enters the market. Due to lack of information about outside options, the firm cannot discriminate in this fashion over a positive measure set of workers in equilibrium. We formally include the assumption that the maximum accepted offer is equal across workers here to rule out equilibria which vary only upon a measure zero set of workers.

A4' is a special case of **A4** and states that off path, conditional on learning the wages of co-workers, workers believe they can receive no more than $w_t^{\sup}$ and will not be rejected if they offer $w_t^{\sup}$. In other words, workers believe that even off path the firm plays a time consistent strategy as in **A0**. Although outside our model, such beliefs are potentially reasonable in the presence of equal pay laws.

All of the results in the paper go through under this expanded game if we restrict attention to equilibria satisfying the above conditions. Indeed, all of the results until those in Section 3.4 go through if relax off path beliefs in **A4'** to workers believing that with probability 1 that any offer weakly less than $w_t^{\sup}$ will be accepted. Nevertheless, this relaxed version of **A4'** can create an additional equilibrium outcome in the game with endogeneous firm selection of transparency in which all firm types pool on $\Lambda = 0$. Further details are available from the authors upon request.

B.2 Double Auction Mapping

This section shows the steps transforming Equation 1 into Equation 2. For $\lambda \in [0, \infty)$

$$\begin{aligned}
& \underset{w_i^*}{\text{argmax}} \left(\tfrac{w_i^*}{\rho+\delta+\lambda} + \tfrac{\lambda}{\rho+\delta+\lambda} \tfrac{\mathbb{E}\left(\bar{w}|\bar{w} \geq w_i^*\right)}{\delta+\rho} \right) \left(1 - \bar{F}(w_i^*)\right) + \tfrac{\theta_i}{\delta+\rho} \bar{F}(w_i^*) \\
= & \underset{w_i^*}{\text{argmax}} \left(w_i^* \tfrac{\rho+\delta}{\rho+\delta+\lambda} + \tfrac{\lambda}{\rho+\delta+\lambda} \mathbb{E}\left(\bar{w}|\bar{w} \geq w_i^*\right) \right) \left(1 - \bar{F}(w_i^*)\right) + \theta_i \bar{F}(w_i^*) \\
= & \underset{w_i^*}{\text{argmax}} \left(w_i^* \tfrac{\rho+\delta}{\rho+\delta+\lambda} + \tfrac{\lambda}{\rho+\delta+\lambda} \mathbb{E}\left(\bar{w}|\bar{w} \geq w_i^*\right) - \theta_i \right) \left(1 - \bar{F}(w_i^*)\right) \qquad (11) \\
= & \underset{w_i^*}{\text{argmax}} \left((1-\Lambda) w_i^* + \Lambda \mathbb{E}\left(\bar{w}|\bar{w} \geq w_i^*\right) - \theta_i \right) \left(1 - \bar{F}(w_i^*)\right) \\
= & \underset{w_i^*}{\text{argmax}} \int_{w_i^*}^1 \left((1-\Lambda) w_i^* + \Lambda x - \theta_i \right) \bar{f}(x) dx
\end{aligned}$$

where $\Lambda = \frac{\lambda}{\rho+\delta+\lambda}$. When $\lambda = \infty$, the scheme is equivalent to a posted price in which

massachusetts-law-salary-history.html) and employers often have little information on workers' outside options in online labor markets such as TaskRabbit. Even if firms are able to observe demographic factors associated with high or low outside options (perhaps such as gender), and would optimally set a different maximum wage for these groups, any such strategy would be in violation of the Equal Pay Act of 1963, opening up the firm to litigation. Therefore, the analysis would be unchanged if instead the firm could observe the demographics of workers but could not select separate wage policies for different groups.

$\Lambda = 1$ (workers receive $\bar{w}$ if they remain at the firm). Therefore, for any $\rho, \delta > 0$ there is a bijection between λ and Λ with higher Λ corresponding to more transparency.

C Omitted proofs

Proof of Proposition 1:

It is easy to see that for all times $t \geq 0$ there exists at least one worker earning $\bar{w}$. Therefore, conditional on receiving wage information, every worker will offer the highest wage visible which equals $\bar{w}$. It remains prove that in equilibrium a worker will never renegotiate without the arrival of wage information. Without loss of generality, let a worker enter the market at $t = 0$ and (in an abuse of notation) let w_s denote the wage the worker offers at time s in the absence of arrival of wage information. If the worker does not renegotiate between times $t^{'}$ and $t^{''}$ then let $w_s = w_{t'}$ for all $s \in [t^{'}, t^{''}]$. Therefore, w is a non-decreasing sequence bounded above by 1. Consider a strategy which dictates in the case of no information arrival that $w_{t'} > w_0$ for some $t^{'} > 0$, such that the worker's expected discounted future surplus at any time s is non-decreasing (this latter condition is clearly necessary for any equilibrium strategy). Consider a "compression" strategy which differs only that sequence w is replaced by w^{ϕ} for $\phi > 1$ such that $w^{\phi}_s = w_{\phi \cdot s}$.

If w is increasing at $t = \alpha \cdot s$ then this also increases the worker's expected discounted future surplus by the equilibrium hypothesis. But then increasing w^{α} at $t = s$ does so as well. Then the compression strategy increases the worker's payoff. Proceeding inductively, it is easy to see that taking $\alpha \to \infty$ (i.e. full compression) is the optimal compression strategy for any w. But the expected payoff from such a compression converges to the payoff from setting $w_s = \lim_{t \to \infty} w_t$ for all $s > 0$. Therefore, the optimal strategy sets $w_s = w^*$ for all $s \geq 0$ and some $w^* \in [0, 1]$.

■

Proof of Proposition 2:

Let $\bar{w} = \beta(v)$ and let $w^*_i = \gamma(\theta)$ and assume that a linear equilibrium exists. Workers are hired at initial wages in some range $[a, h]$ where $0 \leq a \leq h \leq 1$. By the linearity hypothesis, it must be the case that

$$\begin{aligned} \bar{w} &= \begin{cases} v & 0 \leq v < a \\ a + \frac{h-a}{1-a}(v-a) & a \leq v \leq 1 \end{cases} \\ w_i^* &= \begin{cases} a + \frac{h-a}{h}\theta_i & 0 \leq \theta_i \leq h \\ \theta_i & h < \theta_i \leq 1 \end{cases} \end{aligned} \tag{12}$$

Furthermore, by definition $\bar{F}(x) = P(\beta(v) \leq x) = F(\beta^{-1}(x))$, and similarly $\bar{G}(x) = G(\gamma^{-1}(x))$. Inverting the functions in Equation 12 and plugging in to the distributions in Equation 7 yields

$$\begin{aligned} \bar{F}(x) &= 1 - \left(1 - a + \tfrac{(x-a)(1-a)}{h-a}\right)^r \\ \bar{G}(x) &= \left(\tfrac{(x-a)h}{h-a}\right)^s \end{aligned} \qquad a \leq x \leq h \tag{13}$$

Equations 5 and 6 give another set of equations for $\gamma^{-1}(\cdot)$ and $\beta^{-1}(\cdot)$. Plugging these in to the distributions in Equation 7 yields

$$\begin{aligned} \bar{F}(x) = &\ 1 - \left(1 - x - \Lambda\tfrac{\bar{G}(x)}{\bar{g}(x)}\right)^r \\ \bar{G}(x) = &\ \left(x - (1-\Lambda)\tfrac{1-\bar{F}(x)}{\bar{f}(x)}\right)^s \end{aligned} \qquad a \leq x \leq h \tag{14}$$

Solving Equations 13 and 14 simultaneously results in a unique solution in which

$$\begin{aligned} a &= \tfrac{(1-\Lambda)s}{(s+\Lambda)r+(1-\Lambda)s} \\ h &= \tfrac{(1-\Lambda)s+rs}{(s+\Lambda)r+(1-\Lambda)s} \end{aligned} \tag{15}$$

As $\bar{w}$ and w_i^* are pinned down by a and h due to linearity, there is a unique linear equilibrium.

■

Proof of Proposition 3:

The proof of the first two points follows from noting that $\frac{1-\bar{F}(x)}{\bar{f}(x)} = \frac{h-x}{r}$ and $\frac{\bar{g}(x)}{\bar{G}(x)} = \frac{s}{x-a}$ for all $x \in [a, h]$. Both of these terms are strictly decreasing indicating the desired resuilts.

The third and fourth points remain to be shown. We first show $\bar{w}$ is strictly decreasing in Λ for all $v \in [a, 1]$. Using Equations 12 and 15, we see that

$$\bar{w} = a + \tfrac{s}{s+\Lambda}(v-a) \quad \text{for all } v \in [a,1] \tag{16}$$

Differentiating with respect to Λ yields

$$\frac{\partial \bar{w}}{\partial \Lambda} = \frac{\partial a}{\partial \Lambda}\left(1 - \frac{s}{s+\Lambda}\right) - \frac{s}{(s+\Lambda)^2}(v-a) \tag{17}$$

Noting that $\frac{s}{s+\Lambda} \in (0,1]$ and that from Equation 15

$$\frac{\partial a}{\partial \Lambda} \overset{\text{sign}}{=} -r(s+1) < 0 \tag{18}$$

implies that $\frac{\partial \bar{w}}{\partial \Lambda} < 0$ for all $v \in [a,1]$. From Equation 13 we see that $\frac{\bar{G}(x)}{\bar{g}(x)} = \frac{x-a}{s}$ for all $x \in [a,h]$. Therefore, from Equation 6 we see that $\bar{w} \to v$ for all $v \in [0,1]$ as $\Lambda \to 0$.

By virtue of the fact that $\bar{w}$ is decreasing in Λ, it must also be the case that h is decreasing in Λ. (It is possible to directly verify this by computing $\frac{\partial h}{\partial \Lambda}$.) From Equation 13 we calculate $\frac{1-\bar{F}(x)}{\bar{f}(x)} = \frac{h-x}{r}$ for all $x \in [a,h]$. Since h is decreasing in Λ, $\frac{1-\bar{F}(x)}{\bar{f}(x)}$ is also decreasing in Λ over this range. Therefore, from Equation 5 we see that w_i^* is strictly decreasing for $\theta_i \in [0,h]$, and $w_i^* \to \theta_i$ for all $\theta_i \in [0,1]$ as $\Lambda \to 1$.

■

Proof of Theorem 1:

1. For all $\theta_k < h$ we have from Equation 12 that $w_k^* = a + \frac{h-a}{h}\theta_k$. Therefore, for any relevant workers i and j, we have that $w_i^* - w_k^* = \frac{h-a}{h}(\theta_i - \theta_j)$. From Equation 15 we see that the derivative of this function is increasing in Λ, completing the claim.

2. Recall from Equation 11 that the expected lifetime earnings of a worker with outside option θ_i is $T(\Lambda, v, \theta_i) = (1-\Lambda)\, w_i^* + \Lambda \bar{w} - \theta_i$. A sufficient condition for $T(\cdot, v, \theta_i) - T(\cdot, v, \theta_i)$ being strictly decreasing in Λ is that $\frac{\partial^2 T(\Lambda,v,\theta)}{\partial \theta \partial \Lambda} < 0$ for all $\Lambda, \theta \in [0,1)$ and all $v \in [0,1]$. From Equations 11 and 12 we see that

$$\frac{\partial^2 T(\Lambda, v, \theta)}{\partial \theta \partial \Lambda} = \frac{\partial (1-\Lambda)\frac{h-a}{h}}{\partial \Lambda} \tag{19}$$

From Equation 15 we see that

$$\frac{\partial (1-\Lambda)\frac{h-a}{h}}{\partial \Lambda} = \frac{\partial (1-\Lambda)\frac{r}{r+(1-\Lambda)}}{\partial \Lambda} = \frac{-r}{r+1-\Lambda} \tag{20}$$

Since $\Lambda, r > 0$ we see that $\frac{\partial^2 T(\Lambda,v,\theta)}{\partial\theta\partial\Lambda} < 0$ as desired. To show that $T(\cdot, v, \theta_i) - T(\cdot, v, \theta_i) \to 0$ in Λ, we note that $T(\cdot, v, \theta_i) = (1-\Lambda)\, w_i^* + \Lambda\bar{w}$. Since w_i^* is bounded below by θ_i then $T(\cdot, v, \theta_i)$ converges to $\bar{w}(\Lambda)$ for any θ_i.

■

Proof of Theorem 2:

1. To see the equilibrium employment level of the firm, we calculate the probability that a worker is hired by the firm ex-ante. Let $\mathbb{E}(r, s, \Lambda)$ be the expected equilibrium employment level in a market with distribution parameters r and s and transparency Λ. Then

$$\begin{aligned}
\mathbb{E}(r,s,\Lambda) &\equiv \int_0^h Pr\left(\bar{w} \geq w_i^*(\theta)\right) g(\theta) d\theta \\
&= \int_0^h Pr\left(v \geq a + \tfrac{1-a}{h}\theta\right) g(\theta) d\theta \\
&= s \cdot (1-a)^r \int_0^h \left(1 + \tfrac{1}{h}\theta\right)^r \theta^{s-1} d\theta \\
&= (1-a)^r\, h^s \tfrac{\Gamma(r+1)\Gamma(s+1)}{\Gamma(r+s+1)}
\end{aligned} \tag{21}$$

where the first equality comes from substituting in Equation 12, the second equality comes from substituting in the distribution of outside options from Equation 7 and the third from $\Gamma(x) \equiv \int_0^\infty y^{x-1} e^{-y} dy$. As we see, transparency affects employment through changing a and h. We know from Equation 21 that

$$\underset{\Lambda}{\operatorname{argmax}}\, \mathbb{E}(r,s,\Lambda) = \underset{\Lambda}{\operatorname{argmax}}\, (1-a)^r\, h^s \tag{22}$$

Substituting in from Equation 15 and taking the first order condition with respect to Λ yields

$$\Lambda^* = \frac{r+1}{r+s+2} \tag{23}$$

It remains to show that the maximization problem in Equation 22 is concave in Λ over [0,1]. Taking the first order condition of Equation 22 we see that

$$\frac{\partial (1-a)^r h^s}{\partial \Lambda} = -\frac{r^2 s^2 (1-a)^{r-1} h^{r-1} (r(\Lambda-1) + (2+s)\Lambda - 1)}{(s(1+r-\Lambda) + r\Lambda)^3} \tag{24}$$

From this, since $r, s > 0$ and $a < 1$ we see that the first order condition in Equation 23 holds. Substituting in from Equation 7 gives us the particular form of Λ^* in the theorem. We further can calculate

$$\begin{aligned}\frac{\partial^2(1-a)^r h^s}{\partial \Lambda^2} \overset{\text{sign}}{=} & -rsh^s(1-a)^r\left(s^3(r^2+r\left(2-\Lambda^2\right)+\left(1-\Lambda^2\right)\right)\\ & -r\Lambda\left(r^2(2-\Lambda)+2r\left(\Lambda^2-3\Lambda+2\right)+\left(4\Lambda^2-5\Lambda+2\right)\right)\\ & -s^2\left(r^3+r^2\left(-2\Lambda^2+2\Lambda+2\right)+r\left(-2\Lambda^2+4\Lambda+1\right)+2\Lambda\left(1-\Lambda^2\right)\right)\\ & -s\left(r^3\left(-\Lambda^2+2\Lambda+1\right)+r^2\left(3-2\Lambda^2\right)\right)\\ & -s\left(r(6\Lambda^2-6\Lambda+3)+\left(-4\Lambda^3+7\Lambda^2-4\Lambda+1\right)\right)\end{aligned}$$

A sufficient condition for $\frac{\partial^2(1-a)^r h^s}{\partial \Lambda^2} < 0$ for all $\Lambda \in (0,1)$ is that each of the polynomial terms involving Λ be strictly positive for $\Lambda \in (0,1)$. It is easy to check each of these polynomials separately to see that this sufficient condition is indeed satisfied. Therefore, extreme point Λ^* is the global maximizer of expected employment.

2. In equilibrium, there is an outside option cutoff for employment θ^* such that all workers with outside options weakly less than θ^* negotiate wages that are acceptable to the firm. Then employment is equal to $G(\theta^*)$. Noting that a worker i with outside option θ^* sets $w_i^* = \bar{w}$ it must be the case that $G(\theta^*) = \bar{G}(\bar{w})$. We show that $\bar{G}(\bar{w})$ is submodular in v and Λ, and rely on monotone comparative statics techniques of Topkis (1998) to complete the result. From Equations 12 and 13 it is the case that for all $v \geq a$

$$\bar{G}(\bar{w}) = \left(\frac{h}{1-a}(v-a)\right)^s \tag{25}$$

We can use a monotonic transformation of $\bar{G}(\bar{w})$ to complete the claim, that is, we show submodularity of $\frac{h}{1-a}(v-a)$ in v and Λ.

$$\frac{\partial \frac{h}{1-a}(v-a)}{\partial v} = \frac{h}{1-a} = \frac{(1-\Lambda)s + rs}{(s+\Lambda)r} \tag{26}$$

Which is clearly decreasing in Λ. Therefore, $\bar{G}(\bar{w})$ is submodular in v and Λ.

■

Proof of Theorem 3:

We show that the expected equilibrium profit of the firm is strictly increasing in Λ. That the expected equilibrium profit of an arbitrary worker is strictly decreasing in Λ follows a similar calculation. We invoke the law of iterated expectations by first finding the firm's profit for a particular draw $v > a$ which we denote by $\pi(v, \Lambda)$.

$$\begin{aligned}
\pi(v,\Lambda) &= \int_a^{\bar{w}} (v - (1-\Lambda)\, y - \Lambda\bar{w})\, \bar{g}(y)dy \\
&= \int_a^{\bar{w}} (v - (1-\Lambda)\, y - \Lambda\bar{w})\, s \left(\frac{h}{h-a}\right)^s (y-a)^{s-1}dy \qquad (27) \\
&= \frac{(\bar{w}-a)^s}{s+1}\left(\frac{h}{h-a}\right)^s (a\,(\Lambda-1) - \bar{w}(\Lambda+s) + sv + v)
\end{aligned}$$

where the second equality comes by using Equation 13. The ex-ante expected profit of the firm can be expressed as $\pi(\Lambda) = \int_a^1 \pi(v,\Lambda) f(v)dv$. A tedious, but straightforward calculation shows that $\frac{\partial \pi(\Lambda)}{\partial \Lambda} > 0$ for all $r, s > 0$ as desired.

■

Proof of Theorem 4:

We begin this proof with a lemma.

LEMMA **1.** *Let M denote the worker belief that the firm will select $\Lambda = 1$ in equilibrium. Then if $M < 1$ and $\Lambda = c$ and $\Lambda = d$ with $c < d$ are each chosen with positive probability (density) then with positive probability a positive mass of workers will renegotiate wages before learning $\bar{w}$, unless $c = 0$ and $d = 1$.*

Proof of lemma: We assume c and d are chosen with positive probability, but the proof is similar if these are instead selected with positive densities.

$0 \leq c < d < 1$: We prove this case by contraposition. Suppose the set of workers who renegotiate wages before learning $\bar{w}$ has zero measure. Then playing $\Lambda = d$ gives a strictly lower profit than $\Lambda = c$, as it does not change initial bids and only increases the rate at which the firm must pay workers $\bar{w}$. Therefore, any firm setting $\Lambda = d$ has a profitable deviation instead playing $\Lambda = c$, meaning that d cannot be played in equilibrium.

$0 < c < d = 1$: By the previous case we know that there is no $e \in [0, 1) \setminus \{c\}$ that is chosen with positive probability. Therefore, if $\Lambda = 1$ almost every worker will receive $\bar{w}$ in

the instant they arrive at the firm (and therefore never renegotiate) and if $\Lambda = c$ every worker will believe with probability 1 that $\Lambda = c$. Therefore, when $\Lambda = c$ is chosen worker beliefs do not drift over time so by Proposition 1 each worker will again negotiate in the instant they arrive and the instant in which they learn $\bar{w}$. But then any firm setting $\Lambda = c$ has a profitable deviation of instead playing $\Lambda = 0$, meaning that c cannot be played in equilibrium.

■

To complete the proof of the theorem, suppose for contradiction that there is an equilibrium in which $\Lambda = 0$ and $\Lambda = 1$ are played with probabilities p_0 and p_1 where $p_0 > 0$. Then let v_L be the infemum of the firm types that selects $\Lambda = 0$ with positive probability. When a firm plays $\Lambda = 0$ almost all workers (except the zero measure set who learns $\bar{w}$ in the instant they are hired) correctly deduce the firm choice, and that $v \geq v_L$. Since $v = \bar{w}$ when $\Lambda = 0$, each worker will set $w_i^* \geq v_L$. Take a sequence $\{v_\ell\}_{\ell \in \mathbb{N}} \to v_L$ from above. Then for any $\epsilon > 0$ there exists some ℓ^* such that for all $\ell > \ell^*$, $v_\ell - v_L < \epsilon$. Then any firm type v_ℓ with $\ell > \ell^*$ will receive strictly less than ϵ profit per worker it employs when it follows the equilibrium prescription and sets $\Lambda = 0$. Letting $\epsilon \to 0$, any of these firm types with $v_\ell - v_L < \epsilon$ can make higher total profits by selecting $\Lambda = 1$ and setting $\bar{w} = \frac{v}{2}$. Therefore, there can be no such equilibrium in which $p_0 > 0$. By inspection, we have exhausted all cases except that in which $p_1 = 1$. To see that an equilibrium exists in which no worker renegotiates wages and all firms select $\Lambda = 1$, consider the worker belief that $\Lambda = 0$ and $v = 1$ upon not seeing the wages of co-workers at the instant they are hired. The optimal response to these beliefs is to set $w_i^* = 1$ for all i, meaning that the firm will make zero profits if it deviates.

■

Proof of Proposition 4:

Suppressing time and worker indices, suppose a worker has negotiated a flow wage of w. Then in addition to her other choices, she must choose e to solve $\max_{e \in [0,1]} w \cdot e - \theta \cdot e$. For any $w \geq \theta$ the maximizer is $e = 1$.[45] Therefore, when $w \geq \theta$ the equilibrium flow utility to the worker is $w - \theta$, as in the initial model. But by **A1** a worker would never agree to a wage $w < \theta$. So in equilibrium, $e = 1$ and payoffs are the same as the original model. It is easy to see that given this, all other equilibrium choices will be unchanged.

■

[45] Of course, if $w = \theta$ any $e \in [0, 1]$ is a maximizer. For our purposes, we select $e = 1$ in this case, although, as we see, in equilibrium this will only affect a zero measure set of workers.

Proof of Proposition 5:

We prove the second part of the proposition as the first is easy to see.

$\Rightarrow$ Clearly it cannot be the case that $m(e,d) = 0$ for some $d > 0$, or else the firm would never fully equalize wages. Suppose for contradiction that $w \cdot e - \theta \cdot e - m(e,d) = \epsilon > 0$ for some e, d. Let $e^*(d, w_i^*, \bar{w})$ be the optimal effort selected by worker i upon learning $\bar{w}$ and receiving wage w_i^*. The firm must solve

$$\max_{\bar{w} - w_i^* \leq d \leq 0} \; e^*(d, w_i^*, \bar{w})(v - \bar{w} + d) \tag{28}$$

SF1 and **SF2** imply that for any Λ, v, $\bar{w}$ and w_i^*, $d = 0$ is optimal, inducing $e = 1$. This implies that

$$v - \bar{w} \geq e^*(d, w_i^*, \bar{w})(v - \bar{w} + d) \;\; \forall \Lambda,\, d,\, \bar{w},\, v,\, w_i^* \tag{29}$$

which holds if and only if

$$v - \bar{w} \geq d \cdot \frac{e^*(d, w_i^*, \bar{w})}{1 - e^*(d, w_i^*, \bar{w})} \;\; \forall \Lambda,\, d,\, \bar{w},\, v,\, w_i^* \tag{30}$$

Consider a firm of type $v = 1$ and a worker of type $\theta_i = 0$. By sending $r \to \infty$, $h - a \to 1$. By sending $\Lambda \to 0$, the LHS of Equation 30 converges to 0, while by assumption there exists some d for which the RHS is bounded away from 0. Contradiction.

$\Leftarrow$ This direction is easy to see. If $w \cdot e - \theta \cdot e - m(e,d) \leq 0$ for any $e \in [0,1]$ and any $d \in (0,1]$ then as soon as any worker learns $\bar{w}$ the firm can either choose to increase her wage to $\bar{w}$ and receive flow profits $v - \bar{w} \geq 0$ or receive flow profits of 0 otherwise from the worker who will put in zero effort.

■

Proof of Proposition 6:

Taking the first order condition of $\Lambda_m - \Lambda_f$ with respect to λ yields

$$\frac{\alpha_m}{\alpha_f} = \frac{(\rho + \delta + \alpha_m q \lambda)^2}{(\rho + \delta + \alpha_f q \lambda)^2} \tag{31}$$

The LHS of Equation 31 is constant in λ while the RHS is increasing in λ as $\alpha_m > \alpha_f$. Therefore, there is a unique solution λ_c to this first order equation and thus a unique interior

extreme point. As $\Lambda_m - \Lambda_f > 0$ for all $\lambda \in (0, \infty)$ and it is continuously differentiable over this domain, the fact that $\Lambda_m - \Lambda_f = 0$ for $\lambda \in (0, \infty)$ it must be that λ_c is a maximizer, and that $\Lambda_m - \Lambda_f$ is single-peaked.

■

D Multiple firms

In this section, we embed our analysis of pay transparency into a search model by including multiple firms, and show that many of the insights of the main model carry over to this setting. For tractibility, we study only the cases of full secrecy and full transparency. Let $\mathcal{N} = \{1, 2, ..., N\}$ be the set of firms, each with a value for labor v^n drawn iid from distribution F. As before, workers have outside options drawn iid from distribution G. Workers negotiate with firms in a predetermined order without the possibility of returning to an earlier firm. Without loss of generality, we assume that workers first meet with firm 1, then firm 2, and so on.

If a firm rejects a worker's offer the two are ineligible to match at any point in the future, and the worker (instantly) moves to the next firm in the sequence. Although we do not do so for simplicity of exposition, it is possible to embed a search friction in this formulation without affecting the qualitative findings.[46] A worker whose offer is rejected by firm N becomes unemployed for her duration in the market and consumes her outside option. Workers continue to expire at rate ρ at which time they leave the market. A worker whose offer is accepted by firm $n < N$ is replaced with a worker of identical outside option who moves on to firm $n + 1$ as if her offer had been rejected at firm n.[47]

Each firm n selects a maximum wage it is willing to pay for a worker $\bar{w}^n(v^n) \in [0, 1]$, where the choice of $\bar{w}^n$ is not immediately observed by workers. As before, each worker bargains for wages by making TIOLI offers to firms at any point during her employment, potentially renegotiating infinitely often. Workers who at anytime offer a wage greater than $\bar{w}^n$ to firm $n$are permanently unmatched with the firm. Let W_t^n denote the set of wages

[46] Each time a worker's offer is rejected, we could instead make the worker unable to meet with subsequent firms with probability $\zeta \in (0, 1)$. Similarly to the relation between λ and Λ in the main body of the paper, the equilibrium consequences of this probabilistic search friction are identical to a friction which governs the (average) length of time it takes for a worker to find the next firm; in this context ζ close to 0 corresponds to near-instant discovery of the next firm, while ζ close to 1 corresponds to near-infinite time required to discover the next firm. Including such a search friction does not meaningfully change the remainder of the analysis.

[47] This assumption is made for tractability as this "cloning" greatly simplifies equilibrium characterization in our context, and is frequently adopted in the search literature (see, for example, Burdett and Coles (1999), Bloch and Ryder (2000), and Chade (2006)). This assumption may be even more defensible in a setting like TaskRabbit, in which jobs are short-term, and therefore, we can interpret a "cloned" worker as merely a worker who has completed a given task and is not eligible to re-complete it.

firm n is paying to its employed workers, where $W_0^n = \{\bar{w}^n\}$. We model transparency as a random arrival process; at time t, workers matched to firm n matched workers observe W_t^n according to an independent Poisson arrival process with rate $\lambda \in \{0, \infty\}$, where we take $\lambda = \infty$ to mean that the process arrives whenever a worker first matches with a firm.

The timing of the stage game is as follows at each time $t \geq 0$:

1. **Entry:** New workers enter the market. Initialize $m = 1$, and $\ell_i = 1$ for each new worker.

2. **Search and Bargaining:**

   (a) Unmatched workers match with firm m if $\ell_i = m$.

   (b) Each matched worker i learns W_t^m independently with arrival rate λ.

   (c) Newly entering workers must bargain with the firm and any existing, matched worker can initiate bargaining. Any worker i who engages in bargaining makes a TIOLI offer $w_{i,t}^m \in [0, 1]$ to firm m. If $w_{i,t}^m \leq \bar{w}^m$ then firm m pays i a flow wage $w_{i,t}^m$ until i departs or attempts to renegotiate. If $w_{i,t}^m > \bar{w}^m$ then worker i becomes unmatched.

   (d) For any i such that $w_{i,t}^m > \bar{w}^m$ increase ℓ_i by 1.

   (e) If $m < N$, for all i such that $w_{i,t}^m \leq \bar{w}^m$, create a new worker j with $\theta_j = \theta_i$ and $\ell_j = \ell_i + 1$, increase m by 1 and repeat Step 2.

3. **Exit:** Existing workers depart at rate ρ.

D.1 Equilibrium

We work backward to solve for the unique equilibrium. Workers meeting firm N face the same decision as workers in the base model: they face a firm with value v^N drawn from distribution F and are among an incoming cohort with ouside options determined by distribution G. We know from Equations 5 and 6 that under full secrecy each worker i will offer firm N an initial amount w_i^N solving

$$w_i^N - \theta_i = \frac{1 - F(w_i^N)}{f(w_i^N)}$$

and firm N will set $\bar{w}^N = v^N$. Workers will not attempt to renegotiate. Under full transparency, N will set $\bar{w}^N$ to solve

$$v^N - \bar{w}^N = \frac{G(\bar{w}^N)}{g(\bar{w}^N)}$$

and worker i will be employed at flow wage equal to $\bar{w}^N$ if and only iff $\bar{w}^N \geq \theta_i$. Denote by $\theta_i^{n,\lambda}$ the expected equilibrium lifetime utility (under transparency level λ) of a worker with outside option θ_i immediately upon matching with firm n (before making an offer or learning wages through the transparency process), and denote by $G^{n,\lambda}$ the distribution of $\theta_i^{n,\lambda}$. Then, when facing firm $N-1$, workers face will face the same decision but with θ_i replaced with $\theta_i^{N,\lambda}$, and firm $N-1$ will face the same decision as firm N but with distribution G replaced with $G^{N,\lambda}$. Inducting up toward the first firm, we can characterize the equilibrium actions of agents as the following:

$\lambda = 0$:

Workers:

$$w_i^n - \theta_i^{n+1,0} = \frac{1 - F(w_i^n)}{f(w_i^n)} \text{ for } n < N \tag{32}$$

Firms:

$$v^n = \bar{w}^n \text{ for } n \leq N. \tag{33}$$

$\lambda = \infty$:

Workers:

$$w_i^n = \bar{w}^n \mathbf{1}_{\{\bar{w}^n \geq \theta_i^{n+1,\infty}\}} \text{ for } n < N \tag{34}$$

Firms:

$$v^n - \bar{w}^n = \frac{G^{n+1,0}(\bar{w}^n)}{g^{n+1,0}(\bar{w}^n)} \text{ for } n < N. \tag{35}$$

As θ_i is constant over time, $\theta_i^{\cdot,\lambda}$ is a non-increasing sequence, and strictly decreasing for workers with $\theta_i < 1$. Therefore, $\frac{G^{\cdot,\lambda}}{g^{\cdot,\lambda}}(x)$ is non-increasing in n. In words, workers' outside options, which include the option value of bargaining with future firms, decreases as they move along the sequence of firms. Realizing this, under full transparency, earlier firms accept higher wages to incentivize workers to accept their offers rather than wait to meet future firms. We now provide results that are similar to the theorems in the main text.

PROPOSITION **7.** *The expected average utility of workers is higher in equilibrium with $\lambda = 0$ than $\lambda = \infty$. The expected utility of firms is higher in equilibrium with $\lambda = \infty$ than $\lambda = 0$.*

Proof:

We prove this result for workers, and the converse for firms is similar. By Myerson (1981) the expected utility of any worker who reaches firm N is higher under $\lambda = 0$ than $\lambda = \infty$. Therefore, $\theta_i^{N,0} > \theta_i^{N,\infty}$ for all θ_i. When meeting firm $N-1$, worker θ_i is in expectation better off setting offering $\tilde{w}_i^{N-1}$ solving

$$\tilde{w}_i^{N-1} - \theta_i^{N,\infty} = \frac{1 - F(\tilde{w}_i^{N-1})}{f(\tilde{w}_i^{N-1})} \tag{36}$$

than receiving the equilibrium offer under full transparency by the same Myerson (1981) argument. That worker i is able to offer $\tilde{w}_i^{N-1}$ but instead chooses w_i^{N-1} that solves Equation 32 indicates that worker i is better off in expectation by revealed preference under full secrecy. By induction, we see that worker i is better off at every firm she meets under full secrecy.

■

Below are three analogues of the remaining theorems in the body of the paper. The proofs are omitted as the logic follows the proofs of the main theorems.

PROPOSITION **8.** *When $\lambda = \infty$ there is no wage dispersion between workers at the same firm in equilibrium.*

PROPOSITION **9.** *The ex-post employment maximizing level of transparency is weakly decreasing in v.*

PROPOSITION **10.** *When each firm can select $\lambda \in \{0, \infty\}$ as a function of v there is an essentially unique equilibrium outcome. In equilibrium, each firm selects $\lambda = \infty$ for all $v > 0$.*

E Negotiation extensions

In this section, we discuss alternative bargaining protocols that generate qualitatively similar findings as the TIOLI bargaining scheme studied in the body of the paper. The first two cases consider situations in which workers are not able to rebargain as effectively as in the base model, either by being unable to capture the entire difference between their initial offers and $\bar{w}$, or sometimes being unable to rebargain. There is an injection between the equilibria of these games and the game studied in the body of the paper, in which the additional bargaining friction result in de facto lower levels of transparency. The last extension shifts the bargaining power from the workers to the firm probabilistically, giving the firm the ability to propose wages to a fraction of workers.[48] We show that the equilibrium outcome for

[48]This extension is similar to a modeling choice in Halac (2012) which changes the effective bargaining power of two parties by varying the probability of each agent making a TIOLI offer.

workers receiving wage offers is independent of the level of transparency, and the equilibrium outcome for workers proposing wages is identical in this extended game to that of the original game. Therefore, transparency has the same equilibrium effects in this game, just affecting a smaller portion of the workers.

E.1 Workers can only rebargain for part of surplus

There are a number of possibilities as to why workers may not be able to fully close the gap between their initial wage and $\bar{w}$. This could arise from a game in which workers and firm engage in alternating offer bargaining with disagreement amounts set to w_i^*. It could even occur under a worker TIOLI offer scheme under a "non-Markovian" (i.e. does not satisfy condition A0 in Section B.1) equilibrium in which the firm's strategy is equilibrium is to reject rebargaining offers that request more than a fixed proportion of the difference between a worker's initial bid and $\bar{w}$. Formally, suppose that the firm selects $\bar{w}$ which is the maximum wage it accepts from any worker in the initial period a worker is hired. At any subsequent period, the firm rejects any renegotiation offer strictly greater than $w_i^* + \alpha(\bar{w} - w_i^*)$ where $\alpha < 1$.

PROPOSITION **11.** *The (unique) linear equilibrium of the game in which workers can only rebargain for $\alpha \in [0, 1)$ fraction of the difference between $\bar{w}$ and w_i^* and transparency level $\Lambda < 1$ is equivalent to that of the original game with transparency level $\alpha\Lambda$.*

Proof:

For any α the equilibrium of this game is clearly equivalent to that of the original game when $\lambda = \infty$ ($\Lambda = 1$). When $\lambda < \infty$, following the same logic as the main case, workers negotiate at most twice in equilibrium, once when they are first hired, and once when they learn $\bar{w}$ through the transparency process. Letting $\bar{F}(x) = P(\bar{w} \leq x)$, worker i negotiates at the first moment she is hired to:

$$\underset{w_i^*}{\operatorname{argmax}} \left(\frac{w_i^*}{\rho + \delta + \lambda} + \frac{\lambda}{\rho + \delta + \lambda} \left[w_i^* + \alpha \left(\frac{\mathbb{E}\left(\bar{w} | \bar{w} \geq w_i^*\right)}{\delta + \rho} - w_i^* \right) \right] \right) \left(1 - \bar{F}(w_i^*)\right) + \frac{\theta_i}{\delta + \rho} \bar{F}(w_i^*) \tag{37}$$

where the first term represents the weighted (by λ) expected wage the worker receives if matched with the firm, and the second term represents the lifetime earnings of the worker if she exceeds $\bar{w}$ and instead consumes her outside option for her lifetime.

As before, we modify the objective function without affecting the maximizer, and show that this is equivalent to solving:

$$\begin{aligned} & \underset{w_i^*}{\text{argmax}} \int_{w_i^*}^1 \left((1-\Lambda)\, w_i^* + \Lambda \left(w_i^* + \alpha \left(x - w_i^* \right) \right) - \theta_i \right) \bar{f}(x) dx \\ = \quad & \underset{w_i^*}{\text{argmax}} \int_{w_i^*}^1 \left((1-\alpha\Lambda)\, w_i^* + \alpha\Lambda x - \theta_i \right) \bar{f}(x) dx \end{aligned} \tag{38}$$

where $\Lambda = \frac{\lambda}{\rho+\delta+\lambda}$ for all $\lambda \in [0, \infty)$. In equilibrium, the firm sets $\bar{w}(v)$ to

$$\underset{\bar{w}}{\text{argmax}} \frac{\int_0^{\bar{w}} (v-y)\, \bar{g}(y) dy}{\rho+\delta+\lambda} + \bar{G}(\bar{w}) \frac{\lambda}{\rho+\delta+\lambda} \frac{1}{\rho+\delta} \left(v - \left[\int_0^{\bar{w}} y \bar{g}(y) dy + \alpha \left(\bar{w} - \int_0^{\bar{w}} y \bar{g}(y) dy \right) \right] \right) \tag{39}$$

where $\bar{G}(x) = P(w_i^* \leq x)$. The first term gives the total discounted profits made by the firm before the workers experience an event (seeing the wage profile or perishing) and the second term is the profit made from workers after renegotiating their wages to the maximum allowable level over the rest of their lifetimes in the firm. We can similarly manipulate the objective as with the worker problem:

$$\underset{\bar{w}}{\text{argmax}} \int_0^{\bar{w}} \left(v - (1-\alpha\Lambda)\, y - \alpha\Lambda\bar{w} \right) \bar{g}(y) dy \tag{40}$$

Comparing Equations 38 and 40 to Equations 2 and 4, respectively, completes the proof.

■

The results presented in the body of the paper go through in this setting with minor notational changes. The only significant difference is Theorem 2. When α is sufficiently small, the employment maximizing level of transparency may no longer be in the interior. It is possible to show that there exists $\epsilon > 0$ such that for all $\alpha < \epsilon$ full transparency maximizes expected employment if and only if full transparency yields higher expected employment than full secrecy. Intuitively, when α is small, workers are unable to effectively rebargain, creating the possibility that full transparency (which requires no rebargaining in equilibrium) maximizes employment.

E.2 Workers are probabilistically able to rebargain

Now suppose that each worker is able to rebargain with probability α after the first moment she is matched with the firm. Workers who are able to rebargain can take the same actions as in the standard game, while the $1-\alpha$ fraction of workers who cannot rebargain can take no further strategic actions after specifying w_i^*. Workers do not ex-ante know which type they

are, and only realize their type after they make the initial offer to the firm (simultaneously with acceptance or rejection of offer).

PROPOSITION **12.** *The (unique) linear equilibrium of the game in which each worker can independently rebargain with probability $\alpha \in [0, 1)$ and transparency level $\Lambda < 1$ is equivalent to that of the original game with transparency level $\alpha\Lambda$.*

Proof:

Similar to the proof of Proposition 11.

■

Just as with the extension in Appendix E.1, the results from the body of the paper go through with appropriate modification to Theorem 2.

E.3 Two types of workers, receivers and proposers

For this case, suppose that some known fraction of workers are receivers (we can think of these as workers who are bad at bargaining) who are unable to make wage offers to the firm and receive a wage offer from the firm when they are first hired. If a receiver accept the offer, she is locked in to working at the specified wage until she perish. If she rejects, she is permanently unmatched from the firm. The remaining workers operate as before, and make TIOLI offers (potentially infinitely often) to the firm upon matching. The type of each worker is independent of θ_i, and is known to both the worker and the firm.

PROPOSITION **13.** *The (unique) linear equilibrium of the game in which some fraction of workers and receivers and others are proposers is as follows:*

1. *The firm offers all receivers an initial wage of $\bar{\bar{w}}$ which is the same as $\bar{w}$ in the original game with $\Lambda = 1$,*

2. *The (unique) linear equilibrium outcome for proposers with transparency levelΛ is equivalent to that of the original game with transparency Λ.*

Proof:

1. This point follows immediately from the fact that the worker type (proposer or receiver) is independent of θ_i and that θ_i is privately known by each worker.

2. As $\bar{\bar{w}}$ is the optimal posted price wage, the firm cannot maximize profits if it sets $\bar{w} < \bar{\bar{w}}$. Therefore, when any proposer receives wage information through the transaprency process, in equilibrium she will learn $\bar{w} \geq \bar{\bar{w}}$ and will successfully demand a flow wage of $\bar{w}$. Therefore, a proposer's information is not affected by the presence of receivers, and the unique linear equilibrium choices of firm, $\bar{w}$, and proposer, w_i^*, are unchanged from the base model.

■